arXiv:1602.08719v4 [cs.GT] 20 Dec 2016

Envy-Free Pricing in Multi-unit Markets

Simina Brânzei
Hebrew University of Jerusalem, Israel *

Aris Filos-Ratsikas
University of Oxford, United Kingdom †

Peter Bro Miltersen
Aarhus University, Denmark ‡

Yulong Zeng
Tsinghua University, China §

Abstract

We study the envy-free pricing problem in linear multi-unit markets with budgets, where there is a seller who brings multiple units of a good, while several buyers bring monetary endowments. Our goal is to compute an envy-free (item) price and allocation, i.e. an outcome where all the demands of the buyers are met given their budget constraints, which additionally achieves a desirable objective.

In the *price taking* scenario, where the buyers purchase the optimal bundle given the prices they are facing, we provide a polynomial time algorithm for the problem of computing a welfare maximizing envy-free pricing, and an FPTAS and exact algorithm (which is polynomial for a constant number of types of buyers) for the problem of computing a revenue optimal envy-free pricing .

In the *price making* scenario, where the buyers can strategize, we show a general impossibility of designing strategyproof and efficient mechanisms even with public budgets. On the positive side, we provide an optimal strategyproof mechanism whose approximation ratio is a function of the market share, s^*, which can roughly be understood as the maximum buying power of any individual buyer in the market. When the market is even mildly competitive—i.e. with no buyer having a market share higher than 50%—the approximation ratio of our mechanism is at most 2 for revenue and at most $1/(1-s^*)$ for welfare. Moreover, this mechanism is optimal among all the strategyproof mechanisms for both objectives on competitive markets.

Finally, for price taking buyers with general valuations, we provide fully polynomial time approximation schemes as well as hardness results for both revenue and welfare.

Keywords: *envy-free pricing, market share, revenue, social welfare, strategyproofness, efficiency*

*E-mail: simina.branzei@gmail.com
†E-mail: aris.filos-ratsikas@cs.ox.ac.uk
‡E-mail: bromille@cs.au.dk
§E-mail: 297410725@qq.com

1 Introduction

The question of allocating scarce resources among participants with heterogeneous preferences is one of the most important problems faced throughout the history of human society, starting from millenary versions—such as the division of land—to modern variants—such as allocating computational resources to the users of an organization or selling goods in eBay auctions. A crucial development in answer to this question occurred with the invention of money, i.e. of pricing mechanisms that aggregate information about the supply and demand of the goods in order to facilitate trade. The price mechanism was formalized and studied systematically starting with the 19th century, in the works of Fisher [9], who developed a hydraulic model of an economy, and of Arrow-Debreu [2], who developed the exchange market model.

The basic setting is that of a set of participants that come to the market with their initial endowments and aim to purchase goods in a way that maximizes their utility subject to the initial budget constraints. Market equilibria are outcomes where demand and supply meet, have been shown to exist under mild conditions when the goods are perfectly divisible [33], and satisfy very desirable efficiency properties.

The beautiful general equilibrium theory rests on several idealized conditions that are not always met. In particular, the real world is fraught with instances of allocating indivisible goods, such as a house or a piece of jewelry. Market equilibria can disappear in such scenarios due to discontinuities in the valuations; in these models, existence can be guaranteed when the economic agents have essentially unbounded budgets that are of intrinsic value to them (i.e. monetary endowments so large that the agents can always "cover" their valuations for the other goods in the market). However this is a strong assumption that violates basic models such as the Fisher and exchange market, where bounded endowments are an essential element of the profiles of the agents [18, 2].

Nevertheless, we would like to retain at least partially the crisp predictions of market equilibrium theory in the face of indivisibilities, and one method that has been proposed to accomplish this is relaxing the clearing requirement. For instance, in the case of one seller that brings multiple goods, it can be acceptable that at the end of the trade there are some leftover items in stock. The crucial optimality condition of market equilibria is maintained: each buyer purchases an optimal bundle at the current prices. This notion of equilibrium is known as "envy-free" pricing and refers to the fact that no buyer should "envy" any bundle that it could afford in the market [23]. Envy-free pricing schemes are guaranteed to exist in very general models and the question becomes to compute one that recovers some of the efficiency properties of exact market equilibria.

We study the envy-free pricing problem in one of the most basic scenarios possible, namely multi-unit markets with budgets. There is one seller who comes equipped with m units of some good (e.g. jeans), while the buyers bring their money. The seller has no value for the items, while the buyers value both their money and the goods. The goal is to compute an envy-free pricing that has desirable efficiency properties, such as improved welfare or revenue guarantees. This problem falls under the general umbrella of computing

market equilibria and has been studied in a rich body of literature on envy-free auctions, with several NP-hardness results and approximation algorithms for both the revenue and the welfare objectives originating from it. However, the focus has been more on general auctions [23, 5] or auctions with single-minded or unit-demand bidders [23, 15], while some other papers consider the case of unlimited supply of items [23, 11]. Somewhat surprisingly, the linear multi-unit market was not well understood for the simple setting of item pricing and budgets.

Our analysis is carried out for buyers who are *price takers* and *price makers*, respectively. Price taking is the default assumption in the classical market equilibrium models and stipulates that agents are very simple optimizers: they examine the prices computed by the market maker—who knows the preferences—and purchase freely the best bundle available at those prices.[1]

The market can also be seen as a mechanism for allocating goods among finitely many buyers. However, if the market maker does not have enough information about their preferences, the agents may be able to feign different interests in order to influence the prices and eventually get better allocations. This behavior has been referred to as *price making* and it can lead to the emergence of equilibria that are only efficient with respect to the reported—presumably wrong—utility functions [30, 14]. We would like to compute equilibria with improved efficiency guarantees in this much more challenging scenario as well.

1.1 Model and Results

Our model is a multi-unit market with budgets. The utilities are quasi-linear up to the budget cap, and $-\infty$ if the spending exceeds the budget. We study the question of computing envy-free *item* prices, meaning that every unit of a good has the same price, regardless of who buys it or where it is sold. Item pricing is implicit in the standard market models for divisibile goods, is very simple, and is considered the fairest pricing scheme[2].

The problem we study is the following:

Given a linear multi-unit market, compute envy-free prices and allocations with improved welfare and revenue when the buyers are price takers and price makers, respectively.

For the ***price taking*** scenario we provide the following computational results.

[1] An interpretation that sheds light on this assumption is the following: there are infinitely many infinitesimal agents on the market, with a finite number of types (of valuations). Each finite type aggregates the mass (budget) of all the infinitesimal agents described by that valuation. For discrete resources, this perspective essentially implies that many small agents combine their efforts to purchase one good that they like (and will be used by the entire subgroup). However even beyond this interpretation, the pure optimization problem has been of significant independent interest and for this reason it represents the starting point of our analysis.

[2] There exist studies showing that customers are dissatisfied when prices are perceived to be unfair [1].

Theorem A *For linear multi-unit markets with price taking buyers, there exist*

- *a polynomial time algorithm for computing a welfare maximizing envy-free pricing*
- *an FPTAS and an exact algorithm (which runs in polynomial-time for a constant number of types of buyers) for computing a revenue-maximizing envy-free pricing.*

The proof of Theorem A represents the study of Section 3. Our FPTAS for revenue improves upon the results in [22], which had previously provided a 2-approximation algorithm.

The ***price making*** scenario is much more demanding and we deal with the problem of designing envy-free pricings for the strongest concept of incentive compatibility, namely dominant strategy truthfulness. We will evaluate the efficiency of mechanisms using the notion of *market share*, s^*, which captures the buying power that any individual buyer has in the market.

Theorem B *For linear multi-unit markets with known monetary endowments, it is impossible to find an envy-free pricing mechanism that is both efficient and strategyproof. However, there exists an optimal strategyproof mechanism with an approximation ratio bounded by* 2 *for revenue and* $1/(1-s^*)$ *for welfare in competitive markets, i.e. where market share* s^* *is at most* 50%.

The proof of Theorem B is developed in Section 4. The mechanism we construct is optimal among all the strategyproof mechanisms for both objectives on such competitive markets.

Finally, we solve interesting special cases, such as monotone markets for price making buyers (included only in Appendix B) and markets with general valuations for price taking buyers (Appendix C).

1.2 Related Work

There is a direct connection between our model and the literature on envy-free auctions with item-pricing [23]; the market can be interpreted as an auction where prices are set for the units and the payment of a bidder is the sum of prices for the items it obtains. Item pricing has often been advocated as a fair and desirable pricing scheme [22] and the fact that item pricing is inherent in markets makes them worthy candidates for good auction mechanisms [3]. For our setting in particular, our algorithms can be viewed as mechanisms for maximizing revenue and welfare in item-pricing, envy-free multi-unit auctions with budget constraints.

Closely related to our work is the work in [22] in which the authors study very similar types of auctions, aiming for polynomial-time algorithms for maximizing the revenue attained at the best envy-free price. The authors present a *hierarchy of pricing schemes* including (in terms of increasing "customer experience")

item-pricing, item-pricing with upper and lower allocational constraints, bundle pricing and arbitrary pricing, which is equivalent to per-bidder pricing in traditional auctions. Our setting lies in the lowest and most basic level of this hierarchy. Among the other levels of the hierarchy, one of the recently popular ones is bundle pricing (see also, e.g., [16]).

The main results in [22] include "inter-level" bounds, establishing how well pricing schemes at lower levels can approximate the less restricted pricing schemes of the higher levels. As an interesting "intra-level" result, the authors obtain a 2-approximate polynomial-time algorithm for maximizing revenue for item-pricing. We provide an FPTAS which improves upon this 2-approximation.

Many recent works study the envy-free pricing problem through the lens of a different notion, called "pair envy-freeness"[3]; this stipulates that no buyer would want to swap its bundle with that of any other buyer [28, 17]. Monaco, Sankowski, and Zhang [31] considered the item envy-free pricing problem as we do and provide a dynamic programming algorithm for general valuations in multi-unit markets without budgets, which does not imply an FPTAS for our general setting. Chen and Deng [13] also study envy-free pricing in multi-unit markets without budgets and prove hardness results for unit-demand buyers.

The effects of price-making and rational buyers in markets have been studied extensively over the past few years [8, 10, 14, 29, 30]. Given the employment of markets as mechanisms for selling goods, it is particularly important to predict or prevent the consequences of the buyers' selfish behavior. Our paper introduces the first results on strategyproof mechanisms for linear multi-unit markets with budgets. For more general envy-free markets, there has been some work on strategyproof mechanisms in the literature of envy-free auctions [23] and [24] for pair envy-freeness.

Attempts at good strategyproof mechanisms for multi-unit markets are seemingly impaired by a general impossibility result by Dobzinski, Lavi, and Nisan [18] who prove that strategyproofness and efficiency are essentially incompatible when the budgets are *private*. Our general impossibility result is very similar in nature, but it is not implied by the results in [18] for the following two reasons: (a) our impossibility holds for *public* budgets and hence their general theorem does not imply ours and (b) our notion of efficiency is weaker, as it is naturally defined with respect to envy-free allocations only, contrary to the efficiency notion in [18], where envy-freeness is not a constraint. This also means that our impossibility theorem is not implied by their uniqueness result, even for two buyers. A large body of literature considers truthful mechanisms for multi-unit auctions without the envy-freeness constraint (see, e.g., [20, 7, 6, 19]).

Despite a general impossibility, we prove that it is possible to design good strategyproof mechanisms for our objectives, when the monetary endowments are known. Our results for price-making buyers in particular imply that at least in the fundamental setting of envy-free, linear multi-unit markets, strategyproofness does not require discriminatory prices.

[3]Pair envy-freeness is the classical notion of envy-freeness in fair division, and it is implied by the notion used in this paper.

2 Preliminaries

A *linear multi-unit market* is given by a set $N = \{1, \ldots, n\}$ of buyers and a seller that brings m (indivisible) units of a good. Each buyer i is endowed with a budget $B_i > 0$, which has intrinsic value to i, and a valuation $v_i > 0$, which indicates i's value for one unit. The seller has no value for the good and its goal is to extract money from the buyers by selling it, while the buyers aim to maximize their utility by purchasing units of the good as long as the price is profitable.

An *allocation* $\mathbf{x} = (x_1, \ldots, x_n) \in \mathbb{Z}_+^n$ is a vector such that x_i represents the number of units received by buyer i; the allocation must be feasible, i.e. $\sum_{i=1}^n x_i \leq m$.

Given number of units m and a tuple $(\mathbf{v}, \mathbf{B})$, where $\mathbf{v} = (v_1, v_2, \ldots, v_n)$ are the valuation profiles and $\mathbf{B} = (B_1, \ldots, B_n)$ the budgets, an outcome of the market is a pair $(\mathbf{x}, p)$ consisting of an allocation $\mathbf{x}$ and unit-price p.

The *utility* of buyer i at allocation $\mathbf{x}$ and price p is:

$$u_i(p, x_i) = \begin{cases} v_i \cdot x_i - p \cdot x_i, & \text{if } p \cdot x_i \leq B_i \\ -\infty, & \text{otherwise} \end{cases}$$

The *demand* of buyer i at a price p is defined as

$$D_i(p) = \begin{cases} \lfloor \frac{B_i}{p} \rfloor, & \text{if } p < v_i \\ \left\{0, \ldots, \lfloor \frac{B_i}{p} \rfloor\right\}, & \text{if } p = v_i \\ 0, & \text{otherwise.} \end{cases}$$

In other words, the demand set of a buyer is a set consisting of allocations that make the buyer maximally satisfied at the chosen price and respect its budget constraint. Buyer i is said to be *hungry* at price p if $v_i > p$ and *semi-hungry*[4] if $v_i = p$.

Envy-Free Pricing

An allocation and price $(\mathbf{x}, p)$ represent an *envy-free pricing scheme* if $x_i \in D_i(p)$ for all $i \in N$, i.e. each buyer gets a number of units in its demand set at price p. A price p is an *envy-free price* if there exists an allocation $\mathbf{x}$ such that $(\mathbf{x}, p)$ is an envy-free pricing scheme. Similarly, allocation $\mathbf{x}$ *can be supported* at price p if there exists an envy-free tuple $(\mathbf{x}, p)$.

[4]Such buyers are also known as value-constrained [22]. Note that in [22], the demand of semi-hungry buyers is defined to be a singleton consisting of $\lfloor B_i/p \rfloor$, but the definition of envy-freeness does not require that these buyers receive their demand at price p. The two definitions are equivalent and hence all of our results hold for the alternative definition as well.

While an envy-free pricing scheme always exists—just set $p = \infty$—it is not always possible to sell *all* the units in an envy-free way. We illustrate this phenomenon through an example.

Example 1 (Non-existence of envy-free prices that clear the items) *Let $N = \{1, 2\}$, $m = 3$, valuations $v_1 = v_2 = 1.1$, and $B_1 = B_2 = 1$. At any price $p > 0.5$, no more than 2 units can be sold in total because of budget constraints. At $p \leq 0.5$, both buyers are hungry and demand at least 2 units each, but there are only 3 units in total. This example can be generalized to any number of buyers.*

Social Welfare and Revenue

Our main objective is to compute envy-free pricing schemes that have good social welfare and/or revenue. The *social welfare* at an envy-free pricing scheme $(\mathbf{x}, p)$ is

$$\mathcal{SW}(\mathbf{x}, p) = \sum_{i=1}^{n} v_i \cdot x_i$$

while the *revenue* is

$$\mathcal{REV}(\mathbf{x}, p) = \sum_{i=1}^{n} x_i \cdot p$$

3 Markets with Price Taking Buyers

We begin with a few important observations. Because our goal is to maximize the welfare and revenue objective functions, it suffices to only consider *maximal allocations*, i.e. allocations in which one sells as many units as the budgets allow to. Thus throughout sections 3.1 and 3.2) we assume w.l.o.g. that all allocations are maximal.

The optimal solution, for both revenue and welfare, can be found in a set of *candidate prices*: $\mathcal{P} = \left\{ v_i, \frac{B_i}{k} \mid \forall i \in N, \forall k \in [m] \right\}$. These prices are either equal to some valuation or have the property that some buyer could exhaust its budget by purchasing all the units that it can afford.

Crucially, a polynomial time algorithm cannot iterate over the entire set $\mathcal{P}$, because m is given as an integer; the runtime must be polynomial in $\log m$, not m. Finally, w.l.o.g., the buyers are ordered by their valuations: $v_1 \geq v_2 \geq \ldots \geq v_n$.

3.1 Social Welfare in Markets with Price Taking Buyers

Our first main result is a polynomial-time algorithm for computing a welfare-maximizing envy-free pricing scheme.

Theorem 1 *Given a linear multi-unit market, a welfare maximizing envy-free pricing scheme can be computed in polynomial time.*

We first establish a few lemmas.

Lemma 1 *Let p be an envy-free price. Then any price $p' > p$ is also envy-free. Similarly, if p is not an envy-free price, then any price $p' < p$ is not envy-free either.*

Proof: This follows from the fact that for every buyer i, the number of demanded units is non-increasing in the price. If at price p there are enough units to satisfy all demands, then the same holds at any price $p' > p$. Similarly, if at some price p there are not enough units to satisfy all demands, this is also the case for any $p' < p$. □

Lemma 2 *There always exists a social welfare maximizing envy-free price in the set of candidate prices $\mathcal{P}$.*

Proof: Let p' be a welfare maximizing envy-free price and let $\mathbf{x}$ be the corresponding allocation. If $p' \in \mathcal{P}$ then we are doen. Else, assume $p' \neq \mathcal{P}$. Then, we can increase the price p' until some budget B_i is exhausted or the price becomes equal to some valuation v_i. Until that happens, the demand sets of all buyers remain constant and hence the exact same allocation $\mathbf{x}$ can be supported at some price $p \in \mathcal{P}$. Since the social welfare only depends on the allocation and not the price, the conclusion follows. □

Lemma 3 *Let p_{min} be the minimum envy-free price. Then on the interval $[p_{min}, \infty)$, the social welfare is weakly decreasing in the price.*

Proof: Let $\mathbf{x}$ be an envy-free allocation supported at p_{min}. Consider some other price $p' > p_{min}$ and let $\mathbf{x}'$ be a corresponding envy-free allocation supported at p'. Then for each buyer i, the demand weakly decreases from p_{min} to p', and so $x_i \leq x'_i$. Then $v_i \cdot x_i \leq v_i \cdot x'_i$. By summing over all the buyers, we get that $\mathcal{SW}(p_{min}, \mathbf{x}) = \sum_{i=1}^{n} v_i \cdot x_i \leq \sum_{i=1}^{n} v_i \cdot x'_i = \mathcal{SW}(p', \mathbf{x}')$, which completes the proof. □

We can now prove the main theorem.

Proof: (of Theorem 1) We claim that Algorithm 1 computes such a price and allocation. By Lemma 3, to find a welfare maximizing envy-free price it suffices to identify the minimum envy-free price. In fact, as we will explain, it suffices to actually identify a price that is "close enough" to the minimum envy-free price. The algorithm proceeds as follows.

1. First, find an interval of consecutive valuations $[v_{\ell+1}, v_\ell]$ such that $\ell \in N$ is the smallest index for which v_ℓ is an envy-free price while $v_{\ell+1}$ is not (Line 1).

Input: Market with buyers N and m units, linear valuations $\mathbf{v}$, budgets $\mathbf{b}$

Output: Welfare maximizing envy-free pricing $(p^*, \mathbf{x}^*)$

1: Find buyer $\ell \in N$ for which $[v_{\ell+1}, v_\ell]$ contains the welfare maximizing envy-free price
2: Compute $\Delta > 0$ such that $|\frac{B_i}{k} - \frac{B_j}{t}| < \Delta, \forall i, j \in N, \forall k, t \in [m]$
3: $p^* \leftarrow$ SEARCH$(v_{\ell+1}, v_\ell, \Delta)$
4: $\mathbf{x}^* \leftarrow$ COMPUTE-EF-ALLOCATION(p^*) ▷ Given a price, an EF allocation can be computed in $O(n)$
5: **return** $(p^*, \mathbf{x}^*)$

SEARCH(a, b, Δ) ▷ call to a recursive subroutine

6: **if** $b - a < \Delta$ **then**
7: **for** $i \in N$ **do**
8: $k \leftarrow \left\lfloor \frac{B_i}{a} \right\rfloor$
9: **if** $k > \frac{B_i}{b}$ **and** HUNGRY-DEMAND $\left(\frac{B_i}{k}\right) \leq m$ **then**
10: **return** $\frac{B_i}{k}$ ▷ If there is a feasible candidate price in $[a, b]$, output it
11: **end if**
12: **end for**
13: **return** b
14: **end if**
15: **if** HUNGRY-DEMAND $\left(\frac{a+b}{2}\right) > m$ **then**
16: **return** SEARCH$(\frac{a+b}{2}, b, \Delta)$
17: **else**
18: **return** SEARCH$(a, \frac{a+b}{2}, \Delta)$
19: **end if**

HUNGRY-DEMAND(p) ▷ function computing the demand of the hungry buyers at price p

20: **return** $\sum_{i=1:v_i>p}^{n} D_i(p)$

Algorithm 1: COMPUTE-WELFARE-MAXIMIZING-EF-PRICING$(N, \mathbf{v}, \mathbf{b}, m)$

2. Second, compute a length $\Delta > 0$ such that the distance between any two candidate prices in the set $\mathcal{P}$ is smaller than Δ (Line 2).

3. Third, recursively look for the optimal price using binary search over the interval $[v_{\ell+1}, v_\ell]$, stopping when the length of the interval examined is smaller than Δ (Line 3). The price found at this step is optimal and together with the corresponding allocation, it gives the welfare-maximizing envy-free pricing. The HUNGRY-DEMAND function computes the demand of the hungry buyers at a given price p. Note that the price p admits an envy-free allocation if and only if the demand of the hungry buyers can be satisfied, that is HUNGRY-DEMAND$(p) \leq m$.

Next, we argue correctness. For the first step, we examine each valuation v_i in decreasing order and check whether the price $p = v_i$ admits an envy-free allocation. The price p admits an envy-free allocation if and only if the demand of the buyers that are *hungry* can be satisfied. Formally, the condition is:

$$\sum_{j=1:v_j>p}^{n} \left\lfloor \frac{B_j}{p} \right\rfloor \leq m$$

By Lemma 1, it suffices to find an interval $[v_{\ell+1}, v_\ell]$ such that $v_{\ell+1}$ is not envy-free, while v_ℓ is envy-free; the minimum envy-free price is guaranteed to lie in this interval.

For the second step, consider the rational representation of each budget $B_i = \frac{P_i}{Q_i}$. Define

$$\Delta := \min_{i,j \in N} \left(\frac{B_i}{m^2 \cdot P_j \cdot Q_i} \right)$$

For the third step, we compute the price using the subroutine SEARCH invoked on the interval $[v_{\ell+1}, v_\ell]$ with minimum length Δ. The SEARCH(a, b, Δ) function maintains an invariant, namely that at each step, the left endpoint, a, of the examined interval is not an envy-free price, while the right endpoint, b, is. Since the interval $[a, b]$ is small enough (i.e. $b - a < \Delta$), we know that there is at most one candidate price $q = B_i/k$ in $\mathcal{P}$ contained in the interval for some $i \in N, k \in [m]$. If there exists such a price q, we check whether it is envy-free and output it if that is the case. Otherwise, we output b. In either case, by Lemma 2, there exists a minimal price $\bar{p} \in \mathcal{P}$ that is optimal with respect to the social welfare.

Moreover, by Lemma 1, any envy-free price $p' < \bar{p}$, where $p' \notin \mathcal{P}$, guarantees the same social welfare as $\bar{p}$. In the first case, we output $q = \bar{p}$, namely the minimal optimal price in $\mathcal{P}$ whereas in the second case, we output $b < \bar{p}$, which is also optimal by the previous argument.

Finally, we bound the runtime. The first and second step (Lines 1 and 2, respectively) can each be done in time $O(n^2)$. For the third step, we repeatedly halve the interval until its length is bound by Δ. If the initial length is $v_\ell - v_{\ell+1}$, the number of halvings is $O(\log(\frac{v_\ell - v_{\ell+1}}{\Delta}))$, which by definition of Δ is polynomial in the input size. The operations on the interval of length at most Δ (Lines 7-10) can be done in time $O(n^2)$. □

3.2 Revenue in Markets with Price Taking Buyers

In this section, we consider the objective of finding a revenue-maximizing envy-free pricing scheme. For this problem, we design a Fully Polynomial Time Approximation Scheme (FPTAS), i.e. an algorithm that runs in time polynomial in $n, \log m$ and $1/\epsilon$ and obtains an $(1 - \epsilon)$-approximation to the maximum revenue, for any $\epsilon > 0$. Additionally, we prove that the problem can be solved exactly when there is a constant number of *types* of buyers, where the type of a buyer is given by its budget and valuation.

Fully Polynomial Time Approximation Scheme

Our main result in this section is the following theorem.

Theorem 2 *Given a linear multi-unit market, there exists an FPTAS for computing a revenue-optimal envy-free pricing scheme.*

The main idea is the following. If the number of units is small, i.e. at most n/ϵ, then we can iterate over the set $\mathcal{P}$ of candidate prices in polynomial time and find the revenue-maximizing envy-free price and then a corresponding allocation. If m is large, then we consider a continuous variant of the market which is exactly the same as a linear multi-unit market, except that the good is viewed as a continuous resource valued uniformly by each player i, at total value $v_i \cdot m$. For this variant, we prove that we can compute a revenue-maximizing price and allocation in polynomial time and then we *round* the continuous solution to obtain an integral assignment, bounding the loss in revenue due to the rounding procedure.

We begin with some lemmas.

Lemma 4 *Given a linear multi-unit market, let p be a revenue-maximizing envy-free price. Then $p \in \mathcal{P}$, where $\mathcal{P}$ is the set of candidate prices.*

Proof: Let x be any allocation supported by price p and assume for contradiction that $p \notin P$, which means that either p is smaller than all valuations v_i for $i \in N$ or that all buyers have some leftover budget at this price. Then, we could increase the price continuously until at some price p', some budget B_i is exhausted or the price p' becomes equal to some value v_i. In the first case, the demand set of each buyer at price p' is the same as its demand set at price p and in the second case, the demand set of each buyer with $v_j \neq v_i$ at price p' is the same as its demand set at price p, and the demand set of each buyer with $v_i = v_j$ (non-hungry buyers) at price p' contains the demand set of the buyer at price p. In any case, (x, p') is still an envy-free pricing scheme and since the allocation is the same and the price is higher, the revenue at p' is larger than at p. Since p is a revenue-maximizing envy-free price, this is a contradiction. □

The next lemma states that if we had knowledge of an envy-free price that maximizes revenue, we could compute a maximal allocation supported at this price in polynomial time. For this reason, it is sufficient to look for a revenue-maximizing envy-free price.

Lemma 5 *Given an envy-free price p, a revenue maximizing allocation at p can be found in polynomial time in n and $\log(m)$.*

Proof: First, given the valuation functions of the hungry buyers, we can compute their demands at price p. Note these demands are singletons and so the allocation for these buyers is uniquely determined. For the non-hungry buyers (if any), we assign the remaining units (if any) in a greedy fashion: Fix an arbitrary order of buyers and assign them units according to that order, until all of them exhaust their budgets or we run out of units. All these operations can be done in polynomial time. □

At this point, we will consider the *continuous linear multi-unit market*, defined in the beginning of the section. Note that in such a market, every hungry buyer must receive exactly a (B_i/p)-fraction of the good, since the allocation does not have to be integer. We will consider the continuous case because it will be useful in the construction of our algorithm. Finding the revenue-maximizing envy-free price in the continuous variant can be done in polynomial time, as the following lemma suggests.

Lemma 6 *A revenue maximizing envy-free price for the continuous linear multi-unit market can be computed in polynomial time. Moreover, the price is equal to one of the valuations.*

Proof: Recall the valuations are ordered such that $v_1 \geq v_2 \geq \ldots \geq v_n$ and let p be a revenue maxinimizing envy-free price for the continuous problem. If $p = v_i$ for some $i \in N$ we are done, so assume that $p \in (v_\ell, v_{\ell+1})$ for some $\ell \in N$. Since we are in the continuous case and p maximizes revenue, the budgets of buyers $1, 2, \ldots, \ell$ are exhausted at price p, while players $\ell_1, \ldots, n$ have zero demand and the revenue is $\sum_{i=1}^{\ell} B_i$. By setting the price to v_ℓ, we can obtain exactly the same revenue, since the set of buyers that purchases any units is still $\{1, 2, \ldots, \ell\}$ and the budgets of all buyers in this set are still exhausted. By Lemma 1, v_ℓ is an envy-free price and by the argument above, it maximizes revenue. □

The next lemmas concern the revenue achieved in the continuous problem and enable connecting the maximum revenue in the continuous problem to the discrete one.

Lemma 7 *The optimal revenue in a continuous linear multi-unit market is at least as high as in the discrete version.*

Proof: Let (p, x) be a price and allocation that maximize revenue for the discrete problem; denote by R the revenue obtained. Set the price p in the continuous version and consider two cases:

Case 1: The demand is (weakly) lower than the supply at p. Then the same revenue R can be obtained in the continuous version.

Case 2: The demand is higher than the supply. Then we can continuously increase the price until the demand can be met as follows. Let p' be the current price in the continuous problem. If $p' \in (v_\ell, v_{\ell+1})$, then by continuously increasing p' the demand decreases continuously. If the demand is never met in the interior of this interval, then when the price reaches v_ℓ, we can continuously decrease the allocation of player ℓ until either reaching an envy-free pricing or making buyer ℓ's allocation zero (the case where there are multiple

players with valuation v_ℓ is handled similarly, by decreasing their allocations in some sequence). In the latter case, since the demand still exceeds the supply, we iterate by increasing the price continuously in the interval $[v_{\ell-1}, v_\ell)$.

From the two cases, the optimal revenue in the discrete problem is no higher than in the continuous problem, which completes the proof. □

Lemma 8 *Given the optimal revenue R for the continuous linear mulit-unit market, setting the price to $p = R/m$ gives an envy-free pricing scheme with the same revenue where the whole resource is sold.*

Proof: By Lemma 6, the optimal revenue in the continuous problem can be computed in polynomial time by inspecting all the valuations. Let R denote the optimal continuous revenue. If none of the prices where the optimal revenue can be obtained support an allocation at which the entire resource is sold, then the price can be decreased continuously (skipping the valuation points as they have been considered before) while maintaining the revenue constant until the entire resource is sold. By continuity (taking the valuation points into account), there exists a revenue maximizing price at which the whole resource is sold.

This problem can be solved in polynomial time by finding the buyer $\ell \in [n]$ with the property that $\sum_{i=1}^{\ell} B_i = R$. Set $p = \frac{R}{m}$ and compute the corresponding allocation, where buyers $\ell + 1, \ldots, n$ don't receive anything. □

We are now ready to present the algorithm and the corresponding theorem that establishes its correctness and its running time.

Input: Buyers N with linear valuations $\mathbf{v}$, budgets $\mathbf{b}$, m units and $\epsilon > 0$.
Output: An envy-free pricing $(p^*, \mathbf{x}^*)$ with revenue at least $(1-\epsilon)\mathcal{R}$, where $\mathcal{R}$ is the maximum revenue.

if $m \leq \frac{n}{\epsilon}$ **then**
  for each price $p \in \mathcal{P}$ **do**
    Check if p is envy-free. If not, continue with the next candidate price.
    $\mathbf{x}_p \leftarrow$ COMPUTE-EF-ALLOCATION(p)
    $\mathcal{R}_p \leftarrow$ Revenue of $(\mathbf{x}_p, p)$.
  end for
  $p^* \leftarrow \operatorname{argmax}_p \mathcal{R}_p$.
  return $(p^*, \mathbf{x}_{p^*})$
else
  $(\mathbf{x}, p) \leftarrow$ COMPUTE-CONTINUOUS$(N, \mathbf{v}, \mathbf{b}, m)$
  For all $\in N$, let $\bar{x}_i = \lfloor x_i \rfloor$. ▷ Round the allocation down to integers
  return $(p, \bar{\mathbf{x}})$.
end if

Algorithm 2: FPTAS REVENUE-MAXIMIZING-EF-PRICING$(N, \mathbf{v}, \mathbf{b}, m)$

COMPUTE-CONTINUOUS$(N, \mathbf{v}, \mathbf{b}, m)$ ▷ function that computes a revenue-maximizing price where everything is sold in the continuous case

for $i = 1$ to n **do**
  Check if the price v_i is envy-free. If not, continue with the next valuation.
  $\mathbf{x}^i \leftarrow$ COMPUTE-EF-ALLOCATION(v_i).
  $\mathcal{R}^i \leftarrow$ Revenue of $(\mathbf{x}^i, v_i)$.
end for
$\mathcal{R} \leftarrow \max_i \mathcal{R}^i$
$p^* \leftarrow \mathcal{R}/m$.
$\mathbf{x}^* \leftarrow$ COMPUTE-EF-ALLOCATION(p^*).
return$(p^*, \mathbf{x}^*)$

COMPUTE-EF-ALLOCATION(p) ▷ function that computes an envy-free allocation at price p

for each buyer $i \in N$ such that $v_i > p$ **do** ▷ Hungry buyers
  $x_i \leftarrow D_i$ ▷ Demand sets are singletons
end for
for each buyer $i \in N$ such that $v_i = p$ **do** ▷ Non-hungry buyers
  if #available units > 0 **then**
    $x_i \leftarrow \min\{\text{\#available units}, D_i\}$
  end if
end for

Proof: (of Theorem 2) We show that Algorithm 2 is such an FPTAS. Let $\epsilon > 0$ and consider two cases, depending on whether the number of units is *small* or *large*.

Case 1: $m \leq \frac{n}{\epsilon}$. In this case we can just iterate over all the possible prices in $\mathcal{P}$ and select the revenue maximizing envy-free price and the corresponding allocation. This step can be done in time $O\left(\frac{n^2}{\epsilon}\right)$.

Case 2: $m > \frac{n}{\epsilon}$. In this case we first solve the continuous linear mulit-unit market optimally by finding the price and allocation $(p, \mathbf{x})$ where the whole resource is sold. This can be done in polynomial time by Lemma 8. Now consider the allocation $\tilde{\mathbf{x}}$, given by $\tilde{x}_i = \lfloor x_i \rfloor$, i.e. the allocation obtained from $\mathbf{x}$ if we round all fractions down to the nearest integer value. Note that $(p, \tilde{\mathbf{x}})$ is an envy-free pricing in the discrete problem, since each buyer receives their demand. We argue that $(p, \tilde{\mathbf{x}})$ approximates the optimal revenue of the discrete instance within a factor of $1 - \epsilon$.

Let OPT^c and OPT^d denote the optimal revenues in the continuous and discrete instances, respectively,

and R^d the revenue obtained by the rounding procedure above. We have the following inequalities:

$$\frac{R^d}{OPT^d} \geq \frac{OPT^c - p \cdot n}{OPT^d} = \frac{p \cdot m - p \cdot n}{OPT^d} \geq \frac{p \cdot m - p \cdot n}{OPT^c} \geq \frac{p \cdot m - p \cdot n}{p \cdot m} = 1 - \frac{n}{m} > 1 - \epsilon$$

The first inequality holds because the rounding procedure only loses at most pn revenue, the second identity holds because all the units are sold at price p, and the third inequality holds by Lemma 7. This step can be completed in time $O(n^2)$.

The total runtime is bounded by the maximum in each case, which is $O\left(\frac{n^2}{\epsilon}\right)$. □

Exact Algorithm

When the number of types of buyers is constant, a revenue-maximizing envy-free pricing scheme can be computed in polynomial time.

Theorem 3 *Given a linear multi-unit market, a revenue-maximizing envy-free pricing scheme can be computed in polynomial time when the number of types of buyers is constant.*

The omitted proof is included in Appendix A. It is based on the connection between the problem of finding a revenue-maximizing envy-free pricing scheme and the following number theoretic problem.

Problem 1 *Given* $\alpha_1, \alpha_2, ..., \alpha_\ell \in \mathbb{Q}$, *output* $x \in \arg\max \sum_{i=1}^{\ell} \frac{\lfloor \alpha_i x \rfloor}{x}$, $x \in [a, b] \cap \mathbb{N}$.

The idea is that Problem 1 is a generalization of the revenue-maximizing problem. For a constant number of types of players, it can be solved as a special case of the Integer Linear Programming Problem (ILP) [Lenstra [27]], which is known to be polynomial-time solvable for constantly many variables.

4 Markets with Price Making Buyers

Having studied the envy-free pricing problem from the point of view of price taking buyers, we now move to the problem of dealing with rational buyers, who may alter their inputs to improve the outcomes. As standard, we will denote by *mechanism* the algorithm that maps pairs of reported valuations and budgets to prices and allocations.

4.1 Impossibility Result

Our main goal will be incentivize the buyers to reveal their true inputs and our first main result will be an impossibility theorem.

Theorem 4 *There is no strategyproof and Pareto-optimal mechanism to compute an envy-free pricing scheme in linear multi-unit markets, even if the budgets are public.*

We start with the necessary definitions.

Definition 1 *A mechanism for envy-free pricing in linear multi-unit markets is* Pareto-optimal *if for any input v, it outputs an envy-free pricing scheme (p, x) such that there is no other envy-free pricing scheme (p', x') such that for all $i \in N$ it holds that $v_i x' - p'x' \geq v_i x - px$ and $\mathcal{REV}(p', x') \geq \mathcal{REV}(p, x)$ with at least one inequality being strict.*

In other words, a Pareto-optimal mechanism outputs an envy-free price and allocation such that no other envy-free price and allocation is at least as good for all agents and strictly better for at least one agent, where an agent is either a buyer or the auctioneer. Note that this is the very same definition used in [18] but restricted to the set of envy-free prices and allocations, which is the natural definition for our setting.

Recall the definition of an essentially hungry buyer. We will say that given a price p and an allocation $\mathbf{x}$, buyer i is *essentially hungry* if it is either hungry or semi-hungry with $x_i = \lfloor B_i/p \rfloor$. In other words, a buyer is essentially hungry if he receives the non-zero largest element in its demand set.

The following definitions concern the allocation to semi-hungry buyers.

Definition 2 *A mechanism is* demand-maximal *if at any envy-free price p, if it is possible to allocate $\lfloor B_i/p \rfloor$ units to every semi-hungry buyer i, it does.*

It is not hard to see that Pareto-optimality implies demand-maximality. We will be looking to characterize the price of demand-maximal, strategyproof mechanisms. However, there are some "trivial" input profiles for which at any envy-free price, no hungry or semi-hungry buyers can afford a single unit and hence the envy-free price can be anything. A buyer i on profile input v is called *irrelevant* if at any envy-free price p on v, it holds that $v_i < p$ or $v_i \geq p$ and $\lfloor B_i/p \rfloor < 1$, i.e. the buyer can not buy even a single unit. A buyer which is not irrelevant is called *relevant*. Given this definition, a *trivial* profile is a profile $\mathbf{v}$ such that all buyers are *irrelevant*. We will exclude trivial profiles from consideration. In fact, the following lemma proves that when $m > n$, such profiles do not even exist.

Lemma 9 *When $m > n$, there are no trivial valuation profiles.*

Proof: Let v be a valuation profile and assume by contradiction that v is trivial. Let p_{min} be the minimum envy-free price and note that at price p_{min}, hungry and semi-hungry buyers can not afford a single unit. Since p_{min} is the minimum envy-free price however, at any price $p < p_{min}$, there will be over-demand; consider such a price $p = p_{min} - \epsilon$, where ϵ is sufficiently small. At price p, the demand of each buyer must be 1 unit, since it was 0 before and ϵ is small enough. Therefore, the total demand at price p is $n < m$,

contradicting the fact that p is not an envy-free price. $\square$

Finally, we define *in-range* mechanisms: A mechanism is called *in-range* if it always outputs an envy-free price in the interval $[0, v_j]$ where v_j is highest valuation among all buyers that are not irrelevant. Note that any Pareto-optimal mechanism is also in-range.

Now we are ready to prove the main lemma needed for the proof of Theorem 4, which provides a necessary condition for the price of strategyproof, demand-maximal, in-range mechanisms.

Lemma 10 *Let M be a strategy proof, demand-maximal, in-range mechanism and let v be any valuation profile that is not trivial. Then M outputs the minimum envy-free price on v.*

Proof: Assume for contradiction that M does not always output the minimum envy-free price. Let $\mathbf{v} = (v_1, \ldots, v_n)$ be any valuation profile that is not trivial and let p_{min} be the minimum envy-free price on $\mathbf{v}$ and p_v be the price output by M; by assumption, it holds that $p_v > p_{min}$. By the assumption that M is in-range, it holds that $v_j \geq p_v$ for some relevant buyer $j \in N$. Let $J = \{j : v_j \geq p_v : j \text{ is allocated a non-zero number of units } \}$ be the set of all relevant buyers with valuations at least as high as the envy-free price chosen by M.

Now, consider an instance $\mathbf{v^1}$ such that $v_i^1 = v_i$ for all buyers $i \in N \backslash \{j_1\}$ and $p_{min} < v_{j_1}^1 < p_v$ for some buyer $j_1 \in J$, i.e. the instance obtained by $\mathbf{v}$ when some buyer $j_1 \in J$ reports a valuation smaller than p_v and larger than p_{min}. Let p_1 be the price output by M on input $\mathbf{v^1}$.

- Assume first that $p_1 < v_{j_1}^1$. In that case, buyer j_1 on input profile $\mathbf{v}$ would have an incentive to misreport its valuation as $v_{j_1}^1$; that would lower the price and since B_{j_1} is fixed, the buyer would receive at least the same amount of units at a lower price. This would contradict the strategyproofness of M.
- Now consider the case when $p_1 = v_{j_1}^1$. Since $p_1 > p_{min}$, which implies that $\lfloor B_{j_1}/p_1 \rfloor \leq \lfloor B_{j_1}/p_{min} \rfloor$ and since it is possible to allocate $\lfloor B_{j_1}/p_{min} \rfloor$ units to buyer j_1, at price p_1 buyer j_1 is semi-hungry and it is possible to receive the largest element in its demand set. Since M is demand-maximal, it must allocate at least $\lfloor B_i/p_v \rfloor$ units to buyer j_1 at a price $p_1 < p_v$, and buyer j_1 increases its utility by manipulating.

From the discussion above, it must hold that $p_1 > v_{j_1}^1$. For the different instance where buyer 1 has deviated, update the set $J := \{j : v_j \geq p_1 : j \text{ is allocated a non-zero number of units } \}$. If $J = \emptyset$, then Mechanism M is not in-range and we have obtained a contradiction. Otherwise, there must exist some other buyer $j_2 \in J$ with valuation higher than p_1.

Now, consider such a buyer $j_2 \in J$ and the instance $\mathbf{v^2}$ such that $v_i^2 = v_i^1$ for all buyers $i \in N \backslash \{j_2\}$ and $v_{j_1}^1 < v_{j_2}^2 < p_1$ for buyer[5] j_2, i.e. the instance obtained from $\mathbf{v^1}$ when some buyer j_2 in J misreports

[5]Choosing $v_{j_1}^1 < v_{j_2}^2$ makes sure that the size of the set S is decreasing in every step.

its value being between $v^1_{j_1}$ and p_1. Again, assuming that $p_2 \leq v^2_{j_2}$ is the price output by M, buyer j_2 on instance $\mathbf{v^1}$ could benefit from deviating and hence, for the same reasons described earlier, it must be the case that $p_2 > v^2_{j_2}$ and $J \neq \emptyset$, where J is redefined as $J = \{j : v_j \geq p_2\}$. Again, if J is empty, then M is not in-range, otherwise, we can repeat the argument again.

By iteratively considering sequences of valuations obtained in this manner, we eventually obtain an instance $\mathbf{v^{k-1}}$ such that $J = \{j_k\}$, i.e. there is only one buyer with a valuation higher than the envy-free price p_{k-1} output by M. Repeating the argument once more will result in a valuation profile $\mathbf{v^k}$ where the price p_k is higher than the reported valuation $v^k_{j_k}$ of buyer j_k and the set J will be empty, contradicting the fact that M is in-range.

Overall, this implies that M is either not strategyproof, not demand-maximal or not in-range, which is a contradiction. □

Using Lemma 10, we can now prove Theorem 4.

Proof: (of Theorem 4) Assume by contradiction that M is a Pareto-optimal strategyproof mechanism. Consider a market with $n = 2$ buyers and $m = 3$ units. Note that the profile can be adapted to work for any number of buyers, by adding buyers with lower valuations to the market.

Let $v_1 = 3.9$ and $v_2 = 3$ and let $B_1 = B_2 = 6$. The mechanism is Pareto-optimal, demand-maximal and in-range, and so by Lemma 10 and the fact that the profile is not trivial, M must always output the minimum envy-free price. The minimum envy-free price on this profile is 3 and the allocation is such that buyer 1 receives two units (since it is hungry) and buyer 2 receives 1 unit (because receiving two units would violate the unit supply constraint and receiving zero units would violate Pareto-optimality).

Now consider the deviation $v'_1 = 2.01$ of buyer 1 and observe that the new minimum envy-free price is at 2.01. The allocation now gives buyer 1 one unit and buyer 2 two units, for the same reasons mentioned above. The utility of buyer 1 before misreporting was $2(3.9-3) = 1.8$ whereas its utility after misreporting is $1(3.9-2.01) = 1.89$ and the buyer benefits from deviating, violating the assumption that M is strategyproof. □

4.2 Optimal Truthful Mechanism

Our second main result here is an optimal strategyproof mechanism, which has a constant approximation ratio if the market is even mildly competitive. A well-known notion for measuring the competitiveness of a market is the *market share*, understood as the percentage of the market accounted for by a specific entity (see, e.g., the definition of market share and unit market share in Farris, Bendle, Pfeifer, Phillip, and Reibstein [21], Chapter 2).

In our model, the maximum purchasing power (i.e. number of units) of any buyer in the market occurs

at the minimum envy-free price, p_{min}. Recall there are many ways of allocating the semi-hungry buyers, so when measuring the purchasing power of an individual buyer we consider the maximum number of units that buyer can receive, taken over the set of all feasible maximal allocations at p_{min}. Let this set be $\mathcal{X}$. Then the market share of buyer i can be defined as

$$s_i = \max_{\mathbf{x} \in \mathcal{X}} \left(\frac{x_i}{\sum_{k=1}^{n} x_k} \right)$$

The *market share* definition follows: $s^* = \max_{i=1}^{n} s_i$. Roughly speaking, a market share $s^* \leq 1/2$ means that a buyer can never purchase more than half of the resources, case in which our bounds are small constants.

Theorem 5 *There exists a strategyproof mechanism for linear multi-unit markets such that if the market is even mildly competitive, its approximation ratio*

- *for revenue is at most* 2*, whenever* $s^* \leq 1/2$*,*
- *for social welfare is at most* $1/(1 - s^*)$*, whenever* $s^* < 1$*,*

where s^* *is the market share. Moreover, this mechanism is optimal among all the strategyproof mechanisms for both objectives whenever* $s^* \leq 1/2$*.*

ALL-OR-NOTHING Mechanism

Now consider the following mechanism.

MECHANISM **ALL-OR-NOTHING:**

Given a market, let p be the minimum envy-free price and x the allocation obtained as follows:

- For every hungry buyer i, set x_i to its demand.
- For every buyer i with $v_i < p$, set $x_i = 0$.
- For every semi-hungry buyer i, set $x_i = \lfloor B_i/p \rfloor$ if possible, otherwise set $x_i = 0$, considering the semi-hungry buyers in lexicographic order.

In other words, the mechanism always outputs the minimum envy-free price[6] but if there are semi-hungry buyers at that price, they get either all the units they can afford at this price or 0, even if there are

[6]We remark here that we can not always compute the minimum envy-free price, as the price might be irrational. To deal with this, we can assume a discretization of the input space where valuations and candidate prices are points on a grid. The discretization can be fine enough to not effect the ratios for all practical purposes; note that all the results on strategyproofness still hold assuming

still available units, after satisfying the demands of the hungry buyers.

Strategyproofness of the ALL-OR-NOTHING Mechanism

We now prove that ALL-OR-NOTHING is strategyproof for public budgets. More formally, let $I_i = (\mathbf{v_i}, B_i)$ be the type of buyer i and $I = (\mathbf{v}, \mathbf{B})$ a pair of valuation profiles and vector of budgets. As usual, a mechanism M is *strategyproof* if $u_i(M(I)) \geq u_i(M(I'_i, I_{-i})), \forall I, I'_i$, where l_{-i} is the report of the other buyers except i and $I'_i = (\mathbf{v'_i}, B'_i)$ is such that $\mathbf{v'_i} \neq \mathbf{v_i}$, or $B'_i \neq B_i$ or both. That is, no buyer can deviate by reporting a different budget or value and improve its utility. Whenever one of the input parameters is known, the mechanism can be redefined to map only valuations or only budgets to pairs of prices and allocations and all the definitions are similar.

Theorem 6 *Mechanism* ALL-OR-NOTHING *is strategyproof when the budgets are public.*

Proof: First, we will prove the following statement. If p is any envy-free price and p' is an envy-free price such that $p \leq p'$ then the utility of any essentially hungry buyer i at price p is at least as large as its utility at price p'. To see this, first consider the case when $p = p'$. Since buyer i was essentially hungry at price p, and p is envy-free, it received the maximal set of items possible by the budget restriction and d_i be that set. Since the price has not changed, the agent can not receive more than d_i and hence its utility is not larger that before. Now consider the case when $p < p'$. Since p is an envy-free price, buyer i receives a maximal d_i in its demand. For a higher price p', its demand will be at most as large as its demand at price p and hence its utility at p' will be at most as large as its utility at p.

Assume now for contradiction that Mechanism ALL-OR-NOTHING is not strategyproof and let i be a deviating buyer who benefits by misreporting its valuation v_i as v'_i at some valuation profile $\mathbf{v} = (v_1, \ldots, v_n)$, for which the minimum envy-free price is p. Let p' be the new minimum envy free price and let $\mathbf{x}$ and $\mathbf{x'}$ be the corresponding allocations at p and p' respectively, according to ALL-OR-NOTHING. Let $\mathbf{v'} = (v'_i, v_{-i})$ be the valuation profile after the deviation.

We start by arguing that the deviating buyer i is essentially hungry. First, assume for contradiction that i is neither hungry nor semi-hungry, which means that $v_i < p$. Clearly, if $p' \geq p$, then buyer i does not receive any units at p' and there is no incentive for manipulation; thus we must have that $p' < p$. This implies that every buyer j such that $x_j > 0$ at price p is hungry at price p' and hence $x'_j > x_j$. Since the demand of all players does not decrease at p', this implies that p' is also an envy-free price on instance $\mathbf{v}$, contradicting the minimality of p.

Next, assume that buyer i is semi-hungry but not essentially hungry, which means that $v_i = p$ and $x_i = 0$, by the allocation of the mechanism. Again, in order for the buyer to benefit, it has to hold that

the grid input structure. The running time of the mechanism depends on the size of the grid; if the smallest grid interval is inversely polynomial in the size of the input, then the mechanism runs in polynomial time.

$p' < p$ and $x'_i > 0$ which implies that $x'_i = \lfloor B_i/p \rfloor$, i.e. buyer i receives the largest element in its demand set at price p'. But then, since $p' < p$ and p' is an envy-free price, buyer i could receive $\lfloor B_i/p \rfloor$ units at price p without violating the envy-freeness of p. This contradicts the assumption that buyer i is not essentially hungry at price p.

From the previous two paragraphs, the deviating buyer must be essentially hungry. This means that $x_i > 0$ and $v_i \geq p$. By the discussion in the first paragraph of the proof, we have $p' < p$. Since $x_i > 0$, the buyer does not benefit from reporting v'_i such that $v'_i < p'$. Thus it suffices to consider the case when $v'_i \geq p'$. We have two subcases:

- $v'_i > p$: Buyer i is essentially hungry at price p according to v_i and hungry at price p' according to v'_i. The reports of the other buyers are fixed and B_i is public; similarly to above, price p' is an envy-free price on instance $\mathbf{v}$, contradicting the minimality of p.
- $v'_i = p'$: Intuitively, an essentially hungry buyer at price p is misreporting its valuation as being lower trying to achieve an envy-free price p' equal to the reported valuation. Since $v'_i = p'$, Mechanism ALL-OR-NOTHING gives the buyer either as many units as it can afford at this price or zero units. In the first case, since p' is envy-free and B_i is public, buyer i at price p' receives the largest element in its demand set and since the valuations of all other buyers are fixed, p' is also an envy-free price on valuation profile $\mathbf{v}$, contradicting the minimality of p. In the second case, the buyer does not receive any units and hence it does not benefit from misreporting.

Thus there are no improving deviations, which concludes the proof of the theorem. □

Performance of the ALL-OR-NOTHING Mechanism

Next, we show that the mechanism has a good performance for both objectives. We measure the performance of a strategyproof mechanism by the standard notion of approximation ratio, i.e. the worst ratio between the maximum welfare (revenue) over the welfare (revenue) achieved by the mechanism, taken over all valuation profiles:

$$\text{ratio(M)} = \sup_{\mathbf{I}} \frac{\max_{\mathbf{x},\mathbf{p}} \mathcal{OBJ}(I)}{\mathcal{OBJ}(M(I))}$$

where $\mathcal{OBJ} \in \{\mathcal{SW}, \mathcal{REV}\}$ is either the social welfare or the revenue objective. Obviously, a mechanism that outputs a pair that maximizes the objectives has approximation ratio 1. The goal is to construct strategyproof mechanisms with approximation ratio as close to 1 as possible, for either objective.

Recall the definitions of hungry and semi-hungry buyers from the Preliminaries. We will say that given a price p and an allocation $\mathbf{x}$, buyer i is *essentially hungry* if it is either hungry or semi-hungry with $x_i = \lfloor B_i/p \rfloor$. In other words, a buyer is essentially hungry if he receives the non-zero largest element in its demand set.

We say that a valuation profile is *trivial* if at any envy-free price, there is no buyer that can afford to buy a single unit. For the approximation ratios, we only need to consider valuation profiles that are not trivial, since on trivial profiles, both the optimal price and allocation and the price and allocation output by Mechanism ALL-OR-NOTHING obtain zero social welfare or zero revenue.

Theorem 7 *Given a linear multi-unit market, the* ALL-OR-NOTHING *mechanism approximates the optimal revenue within a factor of* 2 *whenever the market share,* s^**, is at most* 50%.

Proof: Let OPT be the optimal revenue, attained at some price p^* and allocation $\mathbf{x}$, and $\mathcal{REV}(AON)$ the revenue attained by the ALL-OR-NOTHING mechanism. By definition, the ALL-OR-NOTHING mechanism outputs the minimum envy-free price p_{min}, together with an allocation $\mathbf{z}$.

There are two cases, depending on whether the optimal envy-free price, p^*, is equal to the minimum envy-free price, p_{min}:

- *Case 1*: $p^* > p_{min}$. Denote by L the set of buyers with valuations at least p^* that can afford at least one unit at the optimal price. Note that the set of buyers that get allocated at p_{min} is a superset of L. Moreover, the optimal revenue is bounded by the revenue attained at the (possibly infeasible) allocation where all the buyers in L get the maximum number of items in their demand.

  These observations lead to the following two inequalities on the revenue attained at the minimum price and the optimal revenue, respectively:

$$\mathcal{REV}(AON) \geq \sum_{i \in L} \left\lfloor \frac{B_i}{p_{min}} \right\rfloor \cdot p_{min}$$

  and

$$OPT \leq \sum_{i \in L} \left\lfloor \frac{B_i}{p^*} \right\rfloor \cdot p^*$$

  Using the previous inequalities, we can bound the revenue of the AON mechanism as follows:

$$\begin{aligned} \frac{\mathcal{REV}(AON)}{OPT} &\geq \frac{\sum_{i \in L} \left\lfloor \frac{B_i}{p_{min}} \right\rfloor \cdot p_{min}}{\sum_{i \in L} \left\lfloor \frac{B_i}{p^*} \right\rfloor \cdot p^*} \\ &\geq \frac{\sum_{i \in L} \left\lfloor \frac{B_i}{p_{min}} \right\rfloor \cdot p_{min}}{\sum_{i \in L} \frac{B_i}{p^*} \cdot p^*} = \frac{\sum_{i \in L} \left\lfloor \frac{B_i}{p_{min}} \right\rfloor \cdot p_{min}}{\sum_{i \in L} B_i} \\ &= \frac{\sum_{i \in L} \left\lfloor \frac{B_i}{p_{min}} \right\rfloor}{\sum_{i \in L} \frac{B_i}{p_{min}}} \geq \frac{\sum_{i \in L} \left\lfloor \frac{B_i}{p_{min}} \right\rfloor}{\sum_{i \in L} 2 \left\lfloor \frac{B_i}{p_{min}} \right\rfloor} = \frac{1}{2}, \end{aligned}$$

where we used the fact that since the market is non-trivial, for any buyer $i \in L$, $\left\lfloor \frac{B_i}{p_{min}} \right\rfloor \geq 1$, and so

$$\frac{B_i}{p_{min}} \leq \left\lfloor \frac{B_i}{p_{min}} \right\rfloor + 1 \leq 2 \left\lfloor \frac{B_i}{p_{min}} \right\rfloor.$$

- *Case 2*: $p^* = p_{min}$. The hungry buyers at p_{min}, as well as the buyers with valuations below p_{min}, receive identical allocations under ALL-OR-NOTHING and the optimal allocation, $\mathbf{x}$. However there are multiple ways of assigning the semi-hungry buyers to achieve an optimal allocation. Recall that $\mathbf{z}$ is the allocation made by ALL-OR-NOTHING. Without loss of generality, we can assume that $\mathbf{x}$ is an optimal allocation with the property that $\mathbf{x}$ is a superset of $\mathbf{z}$ and the following holds:

  (*) *the number of buyers not allocated under* $\mathbf{z}$*, but that are allocated under* $\mathbf{x}$*, is minimized.*

  We argue that $\mathbf{x}$ allocates at most one buyer more compared to $\mathbf{z}$. Assume by contradiction that there are at least two semi-hungry buyers i and j, such that $0 < x_i < \left\lfloor \frac{B_i}{p_{min}} \right\rfloor$ and $0 < x_j < \left\lfloor \frac{B_j}{p_{min}} \right\rfloor$. Then we can progressively take units from buyer j and transfer them to buyer i, until either buyer i receives $x'_i = \left\lfloor \frac{B_i}{p_{min}} \right\rfloor$, or buyer j receives $x'_j = 0$ units.

  Hence we can assume that the set of semi-hungry buyers that receive non-zero, non-maximal allocations in the optimal solution $\mathbf{x}$ is either empty or a singleton.

  If the set is empty, then ALL-OR-NOTHING is optimal. Otherwise, let the singleton be ℓ; denote by $\tilde{x}_\ell$ the maximum number of items that ℓ can receive in any envy-free allocation at p_{min}. Since the number of items allocated by all the maximal envy-free allocations at p_{min} is the same, namely $\sum_{i=1}^n x_i$, but $x_\ell \leq \tilde{x}_\ell$, we get that

  $$\frac{x_\ell}{\sum_{i=1}^n x_i} \leq \frac{\tilde{x}_\ell}{\sum_{i=1}^n x_i} = s_i^*.$$

  We have:

  $$\frac{\mathcal{REV}(AON)}{OPT} = \frac{OPT - x_\ell \cdot p_{min}}{OPT} \geq \frac{OPT - \tilde{x}_\ell \cdot p_{min}}{OPT} = 1 - \frac{\tilde{x}_\ell \cdot p_{min}}{\sum_{i=1}^n x_i \cdot p_{min}}$$
  $$= 1 - \frac{\tilde{x}_\ell}{\sum_{i=1}^n x_i} = 1 - s_i^* \geq 1 - s^*$$

Combining the two cases, it follows that when $s^* \leq 1/2$, the ALL-OR-NOTHING mechanism achieves at least half of the optimum revenue. □

Next we show that the approximation ratio for welfare is also bounded by a constant.

Theorem 8 *The approximation ratio of Mechanism* ALL-OR-NOTHING *with respect to the social welfare is at most* $1/(1 - s^*)$ *in linear multi-unit markets, where the market share* $s^* \in (0, 1)$*. The approximation ratio goes to* 1 *as the market becomes fully competitive.*

Proof: The argument is similar to that in the second case of the revenue analysis (Theorem 7), since the social welfare is maximized at the minimum envy-free price. Specifically, we now have:

$$\frac{\mathcal{SW}(AON)}{OPT} = \frac{OPT - x_\ell \cdot v_\ell}{OPT} \geq \frac{OPT - \tilde{x}_\ell \cdot v_\ell}{OPT} = 1 - \frac{\tilde{x}_\ell \cdot v_\ell}{\sum_{i=1}^n x_i \cdot v_i}$$
$$\geq 1 - \frac{\tilde{x}_\ell \cdot v_\ell}{\sum_{i=1}^n x_i \cdot v_\ell} = 1 - \frac{\tilde{x}_\ell}{\sum_{i=1}^n x_i} = 1 - s_i^* \geq 1 - s^*,$$

where OPT is now the optimal welfare, $\mathbf{x}$ the corresponding allocation at OPT, and we used the fact that $v_\ell \leq v_i$ for all $i \in L$. □

Finally, we will show that Mechanism ALL-OR-NOTHING is optimal among all strategyproof mechanisms for both objectives whenever the market share $s^* \leq 1/2$.

Theorem 9 *Let M be any strategyproof mechanism for linear multi-unit markets. Then the approximation ratio of M for the revenue objective is at least $2 - \frac{4}{m+2}$.*

Proof: Consider a market with equal budgets, B, and valuation profile $\mathbf{v}$. Assume that buyer 1 has the highest valuation, v_1, buyer 2 the second highest valuation v_2, with the property that $v_1 > v_2 + \epsilon$ where ϵ is set later. Let $v_i < v_2$ for all buyers $i = 3, 4, \ldots, n$. Set B and ϵ such that $\lfloor \frac{B}{v_2} \rfloor = \lfloor \frac{B}{v_2+\epsilon} \rfloor = \frac{m}{2} + 1$. Informally, the buyers can afford $\frac{m}{2} + 1$ units at prices v_2 and $v_2 + \epsilon$. Note that on this profile, Mechanism ALL-OR-NOTHING outputs price v_2 and allocates $\frac{m}{2} + 1$ units to buyer 1.

Let M be any strategyproof mechanism, p_M its price on this instance, and p^* the optimal price (with respect to the objective in question). The high level idea of the proof, for both objectives, is the following. We start from the profile $\mathbf{v}$ above, where $p_{min} = v_2$ is the minimum envy-free price, and argue that if $p_M \neq v_2$, then the bound follows. Otherwise, $p_M = v_2$, case in which we construct a sequence of profiles, starting from $\mathbf{v}$ — $\mathbf{v}^{(1)}, \mathbf{v}^{(2)}, \ldots, \mathbf{v}^{(k)}$ — that only differ from the previous profile in the sequence by the reported valuation $v_2^{(j)}$ of buyer 2. We argue that in each such profile, either the mechanism allocates items to buyer 1 only, case in which the bound is immediate, or buyer 2 is semi-hungry. In the latter case, strategyproofness and the constraints on the number of items will imply that any truthful mechanism must allocate to buyer 2 zero items, yielding again the required bound.

First, consider the social welfare objective. Observe that for the optimal price p^* on profile $\mathbf{v}$, it holds that $p^* = v_2$. We have a few subcases:

- *Case 1*: $p_M < v_2$. Then M is not an envy-free mechanism, since in this case there would be over-demand for units.
- *Case 2*: $p_M > v_2$: Then M allocates units only to buyer 1, achieving a social welfare of at most

$(\frac{m}{2}+1)v_2$. The maximum social welfare is $m \cdot v_2$, and so the approximation ratio of M is at least

$$\frac{m}{(m/2)+1} = 2 - \frac{4}{m+2}.$$

- *Case 3*: $p_M = v_2$: Let x_2 be the number of units allocated to buyer 2 at price v_2; note that since buyer 2 is semi-hungry at v_2, any number of units up to $\frac{m}{2}-1$ is a valid allocation. If $x_2 = 0$, then M allocates units only to buyer 1 at price v_2 and for the same reason as in Case 2, the ratio is at least $2 - \frac{4}{m+2}$; therefore we can assume $x_2 \geq 1$.

  Next, consider the valuation profile $\mathbf{v^{(1)}}$ where for each buyer $i \neq 2$, it holds that $v_i^{(1)} = v_i$, while for buyer 2 we have $v_2 < v_2^{(1)} < v_2 + \epsilon$. By the definition of B, the minimum envy-free price on $\mathbf{v^{(1)}}$ is $v_2^{(1)}$. Let $p_M^{(1)}$ be the price output by M on valuation profile $\mathbf{v}^{(1)}$ and consider a few subcases:

  (i) $p_M^{(1)} > v_2^{(1)}$: Then using the same argument as in Case 2, the approximation is at least $2 - \frac{4}{m+2}$, and we are done.

  (ii) $p_M^{(1)} < v_2^{(1)}$: This cannot happen because by definition of the budgets, $v_2^{(1)}$ is the minimum envy-free price.

  (iii) $p_M^{(1)} = v_2^{(1)}$: Let $x_2^{(1)}$ be the number of units allocated to buyer 2 at profile $\mathbf{v^{(1)}}$; we claim that $x_2^{(1)} \geq 2$. Otherwise, if $x_2^{(1)} \leq 1$, then on profile $\mathbf{v^{(1)}}$ buyer 2 would have an incentive to misreport its valuation as v_2, which would move the price to v_2, giving buyer 2 at least as many units (at a lower price), contradicting strategyproofness.

  Next, consider the valuation profile $\mathbf{v^{(2)}}$, where for each buyer $i \neq 2$, it holds that $v_i^{(2)} = v_i^{(1)} = v_i$ and for buyer 2 it holds that $v_2^{(1)} < v_2^{(2)} < v_2 + \epsilon$. For the same reasons as in Cases (i)-(iii), the behavior of M must be such that:

  – the price output on input $\mathbf{v^{(2)}}$ is $v_2^{(2)}$ (otherwise M only allocates to buyer 1, and the bound is immediate) and

  – the number of units $x_2^{(2)}$ allocated to buyer 2 is at least 3 (otherwise strategyproofness would be violated).

  By iterating through all the profiles in the sequence, we arrive at $\mathbf{v^{(k)}}$ (similarly constructed), where the price is $v_2^{(k)}$ and buyer 2 receives at least $m/2$ units. However, buyer 1 is still hungry at price $v_2^{(k)}$ and should receive at least $\frac{m}{2}+1$ units, which violates the unit supply constraint. This implies that in the first profile, $\mathbf{v}$, M must allocate 0 units to buyer 2 (by setting the price to v_2 or to something higher where buyer 2 does not want any units). This implies that the approximation ratio is at least $2 - \frac{4}{m+2}$.

For the revenue objective, the argument is exactly the same, but we need to establish that at any profile $\mathbf{v}$ or $\mathbf{v^{(i)}}$, $i = 1, \ldots, k$ that we construct, the optimal envy-free price is equal to the second highest reported

valuation, i.e. v_2 or $v_2^{(i)}$, $i = 1, \ldots, k$ respectively. To do that, choose v_1 such that $v_1 = v_2 + \delta$, where $\delta > \epsilon$, but small enough such that $\lfloor \frac{B}{v_2+\delta} \rfloor = \lfloor \frac{B}{v_2} \rfloor$, i.e. any hungry buyer at price $v_2 + \delta$ buys the same number of units as it would buy at price v_2. Furthermore, ϵ and δ can be chosen small enough such that $(\frac{m}{2} + 1)(v_2 + \delta) < m \cdot v_2$, i.e. the revenue obtained by selling $\frac{m}{2} + 1$ units to buyer 1 at price $v_2 + \delta$ is smaller than the revenue obtained by selling $\frac{m}{2} + 1$ units to buyer 1 and $\frac{m}{2} - \epsilon$ units to buyer 2 at price v_2. This establishes that the optimal envy-free price is the same as before, for every profile that we consider and all the arguments go through. □

In Section B of the appendix, we show that there are interesting special cases where the approximation ratio of ALL-OR-NOTHING is optimal regardless of the market share.

5 Discussion

To summarize, we have considered the envy-free pricing problem in one of the most basic settings possible, that of item pricing in linear multi-unit markets, where the players have bounded monetary endowments.

For markets with price taking buyers we design a polynomial time algorithm for maximizing social welfare, while for the problem of maximizing revenue we give an FPTAS and an exact algorithm that runs in polynomial time for constantly many types of buyers. Our algorithms for the price taking scenario work more generally for the (multi-unit) Fisher market model, which differs by the fact that the buyers have no intrinsic value for their budgets. The Fisher market model with discrete goods has been studied in the context of the approximate CEEI mechanism, which is a different relaxation of the market equilibrium that involves randomization (see, e.g., [32], [12]). The exact CEEI solution has been studied as well for discrete goods with perfect substitutes and perfect complements in [4].

For price making buyers, we find an impossibility showing that no strategyproof and Pareto efficient mechanism exists, even when the budgets are public. However, we design an optimal truthful mechanism, which simultaneously approximates the optimal revenue and social welfare within a constant factor when the market is even mildly competitive (i.e. has market share at most 50%). These results require that the monetary endowments are known.

We additionally study several interesting cases, such as monotone markets (see Appendix B), as well as markets with *general* valuations (see Appendix C), where the buyers are equipped with a different value for every possible allocation. We find that the problem of computing the optimal welfare or revenue envy-free pricing is NP-hard; however there exists an FPTAS for both objectives, which exploits a connection with the MULTI-CHOICE KNAPSACK problem.

An intriguing question for future work is finding the exact complexity of revenue maximization for envy-free linear multi-unit markets. In Section D of the Appendix, we prove that Problem 1, which allowed us to

establish a polynomial-time algorithm for solving the problem for a constant number of buyers, is actually NP-hard for many buyers and hence unless P=NP, it can not be used to obtain a polynomial-time algorithm for the general case.

Finally, it would be very interesting to obtain a full characterization of strategyproofness in this domain.

6 Acknowledgements

We would like to thank Branislav Bosansky and Kristoffer Arnsfelt Hansen for useful discussion.

Simina Brânzei was supported by (I) ISF grant 1435/14 administered by the Israeli Academy of Sciences and Israel-USA Bi-national Science Foundation (BSF) grant 2014389, (II) the I-CORE Program of the Planning and Budgeting Committee and The Israel Science Foundation, and (III) the Simons Institute for the Theory of Computing. Aris Filos-Ratsikas was also partially supported by the ERC Advanced Grant 321171 (ALGAME). Finally, the authors aknowledge support from the Danish National Research Foundation and The National Science Foundation of China (under the grant 61361136003) for the Sino-Danish Center for the Theory of Interactive Computation and from the Center for Research in Foundations of Electronic Markets (CFEM), supported by the Danish Strategic Research Council. Simina Brânzei and Aris Filos-Ratsikas also received partial support from the COST Action IC1205 on "Computational Social Choice"

References

[1] Eric Anderson and Duncan Simester. Price stickiness and customer antagonism. *Available at SSRN 1273647*, 2008.

[2] Kenneth J. Arrow and Gerard Debreu. Existence of an equilibrium for a competitive economy. *Econometrica*, 22(3):265–290, 1954.

[3] Moshe Babaioff, Brendan Lucier, Noam Nisan, and Renato Paes Leme. On the efficiency of the walrasian mechanism. In *Proceedings of the fifteenth ACM conference on Economics and computation*, pages 783–800. ACM, 2014.

[4] Simina Brânzei, Hadi Hosseini, and Peter Bro Miltersen. Characterization and computation of equilibria for indivisible goods. In *Proceedings of the 8th International Symposium on Algorithmic Game Theory (SAGT)*, pages 244–255. Springer, 2015.

[5] Maria-Florina Balcan, Avrim Blum, and Yishay Mansour. Item pricing for revenue maximization. In *Proceedings of the 9th ACM conference on Electronic commerce*, pages 50–59. ACM, 2008.

[6] Yair Bartal, Rica Gonen, and Noam Nisan. Incentive compatible multi unit combinatorial auctions. In *Proceedings of the 9th conference on Theoretical aspects of rationality and knowledge*, pages 72–87, 2003.

[7] Christian Borgs, Jennifer Chayes, Nicole Immorlica, Mohammad Mahdian, and Amin Saberi. Multi-unit auctions with budget-constrained bidders. In *Proceedings of the 6th ACM Conference on Electronic Commerce*, EC '05, pages 44–51, 2005.

[8] Allan Borodin, Omer Lev, and Tyrone Strangway. Budgetary effects on pricing equilibrium in online markets. *arXiv preprint arXiv:1511.06954*, 2015.

[9] William Brainard and Herbert E. Scarf. How to compute equilibrium prices in 1891. *Cowles Foundation Discussion Paper*, 2000.

[10] Simina Brânzei, Yiling Chen, Xiaotie Deng, Aris Filos-Ratsikas, Søren Kristoffer Stiil Frederiksen, and Jie Zhang. The fisher market game: Equilibrium and welfare. In *Twenty-Eighth Conference on Artificial Intelligence*, pages 587–593, 2014.

[11] Patrick Briest and Piotr Krysta. Single-minded unlimited supply pricing on sparse instances. In *Proceedings of the seventeenth annual ACM-SIAM symposium on Discrete algorithm*, pages 1093–1102. Society for Industrial and Applied Mathematics, 2006.

[12] E. Budish. The combinatorial assignment problem: Approximate competitive equilibrium from equal incomes. *Journal of Political Economy*, 6(119):1061–1103, 2011.

[13] Ning Chen and Xiaotie Deng. Envy-free pricing in multi-item markets. *ACM Transactions on Algorithms (TALG)*, 10(2):7, 2014.

[14] Ning Chen, Xiaotie Deng, and Jie Zhang. How profitable are strategic behaviors in a market? In *ESA*, pages 106–118. Springer, 2011.

[15] Maurice Cheung and Chaitanya Swamy. Approximation algorithms for single-minded envy-free profit-maximization problems with limited supply. In *Foundations of Computer Science, 2008. FOCS'08. IEEE 49th Annual IEEE Symposium on*, pages 35–44. IEEE, 2008.

[16] Edith Cohen, Michal Feldman, Amos Fiat, Haim Kaplan, and Svetlana Olonetsky. Truth, envy, and truthful market clearing bundle pricing. In *Proceedings of the 7th International Conference on Internet and Network Economics (WINE)*, pages 97–108, 2011.

[17] Riccardo Colini-Baldeschi, Stefano Leonardi, Piotr Sankowski, and Qiang Zhang. Revenue maximizing envy-free fixed-price auctions with budgets. In *Web and Internet Economics*, pages 233–246. Springer, 2014.

[18] Shahar Dobzinski, Ron Lavi, and Noam Nisan. Multi-unit auctions with budget limits. *Games and Economic Behavior*, 74(2):486–503, 2012.

[19] Shahar Dobzinski and Noam Nisan. Mechanisms for multi-unit auctions. In *Proceedings of the 8th ACM conference on Electronic commerce*, pages 346–351, 2007.

[20] Shahar Dobzinski and Noam Nisan. Multi-unit auctions: beyond roberts. volume 156, pages 14–44. Elsevier, 2015.

[21] Paul W Farris, Neil T Bendle, Phillip E Pfeifer, and David J Reibstein. *Marketing metrics: The definitive guide to measuring marketing performance*. Pearson Education, 2010.

[22] Michal Feldman, Amos Fiat, Stefano Leonardi, and Piotr Sankowski. Revenue maximizing envy-free multi-unit auctions with budgets. In *Proceedings of the 13th ACM Conference on Electronic Commerce*, pages 532–549. ACM, 2012.

[23] Venkatesan Guruswami, Jason D Hartline, Anna R Karlin, David Kempe, Claire Kenyon, and Frank McSherry. On profit-maximizing envy-free pricing. In *Proceedings of the sixteenth annual ACM-SIAM symposium on Discrete algorithms*, pages 1164–1173. Society for Industrial and Applied Mathematics, 2005.

[24] Jason Hartline and Qiqi Yan. Envy, truth, and profit. In *Proceedings of the 12th ACM Conference on Electronic Commerce (EC)*, pages 243–252, 2011.

[25] Jason D Hartline. Mechanism design and approximation. *Book draft. October*, 122, 2013.

[26] Eugene L. Lawler. Fast approximation algorithms for knapsack problems. *Mathematics of Operations Research*, 4(4):339–356, 1979.

[27] Hendrik W Lenstra Jr. Integer programming with a fixed number of variables. *Mathematics of operations research*, 8(4):538–548, 1983.

[28] Evangelos Markakis and Orestis Telelis. Envy-free revenue approximation for asymmetric buyers with budgets. In *Proceedings of SAGT 2016*, 2016. to appear.

[29] Ruta Mehta and Milind Sohoni. Exchange markets: Strategy meets supply-awareness. In *Web and Internet Economics*, pages 361–362. Springer, 2013.

[30] Ruta Mehta, Nithum Thain, László A Végh, and Adrian Vetta. To save or not to save: The fisher game. In *Web and Internet Economics*, pages 294–307. Springer, 2014.

[31] Gianpiero Monaco, Piotr Sankowski, and Qiang Zhang. Revenue maximization envy-free pricing for homogeneous resources. In *Proceedings of the 24th International Conference on Artificial Intelligence (IJCAI)*, pages 90–96, 2015.

[32] Abraham Othman, Christos H. Papadimitriou, and Aviad Rubinstein. The complexity of fairness through equilibrium. In *Proceedings of the ACM Conference on Economics and Computation (EC)*, pages 209–226, 2014.

[33] Léon Walras. Elements d'economie politique pure, ou theorie de la richesse sociale (in french). 1874. English translation: Elements of pure economics; or, the theory of social wealth. American Economic Association and the Royal Economic Society, 1954.

APPENDIX

A Exact Algorithm

In Section 3.2 of the main text, we presented Theorem 3 which states that when the number of buyers is constant, we can find a revenue-maximizing price and allocation in polynomial time and we gave a high-level idea of the proof. Note that the algorithm also works for a constant number of buyer *types*, which can be proven using very similar arguments. In this section, we present the complete proof of Theorem 3. Before proving the theorem, for ease of exposition, we re-state the problem that generalizes the problem of computing a revenue-maximizing envy-free pricing scheme.

Problem 1 *Given* $\alpha_1, \alpha_2, ..., \alpha_\ell \in \mathbb{Q}$, *output* $x \in \arg\max \sum_{i=1}^{\ell} \frac{\lfloor \alpha_i x \rfloor}{x}$, $x \in [a, b] \cap \mathbb{N}$.

Lemma 11 *If Problem 1 can be solved in polynomial time, then the problem of finding a revenue-maximizing envy-free pricing scheme can be solved in polynomial time as well.*

Proof: First, recall that by Lemma 4, the revenue-maximizing envy-free price p is in set $\mathcal{P}$, i.e. it is either equal to some valuation v_i or equal to some fraction B_i/k, for some $i \in N$ and some $k \in [1, m] \cap \mathbb{N}$, which means that at least one of the budgets is exhausted at that price and allocation. In the former case, we can simply iterate over all n values $v_1, \ldots, v_n$ and the revenue-maximizing price can be found in polynomial time and by Lemma 5, so can the corresponding revenue-maximizing allocation and the lemma holds. Therefore, we will assume that $p = B_i/k$ for some $i \in N$ and some integer $k \leq m$. We will show how to find the revenue-maximizing price given that some buyer's j budget is exhausted; then we can iterate over all the buyers and calculate the revenue obtained in each case, outputting the price and allocation with the maximum revenue.

Let $\alpha_i = B_i/B_j$ and let x be the allocation of buyer j. From the discussion above, it holds that $p = B_j/x$. Assuming that buyer i is hungry at price p, the revenue obtained by buyer i is

$$p \cdot \left\lfloor \frac{B_i}{p} \right\rfloor = \frac{B_j}{x} \cdot \left\lfloor \frac{\alpha_i \cdot B_j}{x} \right\rfloor = B_j \frac{\lfloor \alpha_i x \rfloor}{x},$$

The total revenue is $\mathcal{REV} = B_j \sum_{i=1}^{l} \frac{\lfloor \alpha_i x \rfloor}{x}$, where l is the number of hungry buyers at price $p = B_j/x$.

Now, given that a revenue-maximizing envy-free price is not equal to some valuation v_i, it lies in an interval (v_{i+1}, v_i) for some $i \in N$. For each such interval $\mathcal{I}$, the set of hungry buyers $S_{\mathcal{I}}$ at any chosen price consists of the buyers with valuations $v_j \geq v_i$ and hence we know exactly the members of this set. By the discussion above, by letting $\ell = |S_{\mathcal{I}}|$ in Problem 1, the value of the quantity $\sum_{i=1}^{\ell} \frac{\lfloor a_i x \rfloor}{x}$ gives exactly the maximum revenue attained in the interval (v_{i+1}, v_i). By iterating over all intervals, we can find the revenue-maximizing price and allocation.

What is left to show is that the envy-freeness constraint can be captured by the constraint $x \in [a, b] \cap \mathbb{N}$. Note that the fact that the price lies in some open interval between two valuations and the envy-freeness constraint define an interval $[c, d]$ from which the price has to be chosen. The value of c is either the minimum envy-free price in $\mathcal{P}$ or the smallest price in $\mathcal{P}$ strictly larger than v_{i+1} whereas the value of d is the largest price in $\mathcal{P}$ strictly smaller than v_i. All of these prices can be found in polynomial time by binary search on the set $\mathcal{P}$, similarly to Algorithm 1. Therefore, the price constraint can be written as $p \in [c, d] \cap \mathcal{P}$ and by setting $a = B_j/d$ and $b = B_j/c$, we obtain the corresponding interval of Problem 1. □

By Lemma 11, it is sufficient to prove that Problem 1 can be solved in polynomial time; the same algorithm will also find a revenue-maximizing envy-free pricing scheme. By setting $k_i = \lfloor a_i x \rfloor$, from Problem 1 we obtain the following equivalent problem:

Problem 2 *Given $a, b, \alpha_i \in \mathbb{Q}$,*

$$\begin{aligned} \max \quad & \frac{k_1 + ... + k_\ell}{x} \\ s.t. \quad & k_i \leq \alpha_i x < k_i + 1, \;\; i = 1, 2, ..., \ell \\ & a \leq x \leq b \\ & x, k_i \in \mathbb{Z}, \;\; i = 1, 2, ..., \ell \end{aligned}$$

Problem 2 is obviously an optimization problem with a non-linear objective function. We will consider the decision version of the problem, where the value of the objective function is restricted to lie in some interval.

Problem 3 *Given $a, b, c, d, \alpha_i \in \mathbb{Q}$, decide if there exist $x, k_1, \ldots, k_n \in \mathbb{Z}$ such that*

$$\begin{aligned} & cx \leq k_1 + ... + k_n \leq dx \\ & k_i \leq \alpha_i x < k_i + 1, \;\; i = 1, 2, ..., n \\ & a \leq x \leq b \\ & x, k_i \in \mathbb{Z}, \;\; i = 1, 2, ..., n \end{aligned}$$

We are now ready to prove Theorem 3.

Proof: (of Theorem 3) We will prove that Problem 2 can be solved in polynomial time, by proving that Problem 4 can be written as an Integer Linear Program (ILP) problem and then providing an algorithm for Problem 2 given an algorithm for Problem 4. Since Problem 1 and Problem 2 are equivalent (by Lemma 11), this will also establish the existence of a polynomial time algorithm for finding a revenue-maximizing envy-free pricing scheme. First, the ILP problem is known to be solvable in polynomial time when the number of variables is fixed.

Lemma 12 ([27]) *The ILP problem with a fixed number of variables can be solved in polynomial time.*

In order to show that Problem 4 can be written as an ILP, we need to handle the non-integer coefficients. For every coefficient $j \in \mathbb{Q}$, we multiply their common denominator in each equation and all coefficients become integers, without affecting the length of the input. The inequalities that have the "$\geq$" direction can be handled easily. To handle the strict inequality constraint, for each such constraint such that $\mathcal{A} < \mathcal{B}$, we instead write $2\mathcal{A} \leq 2\mathcal{B} - 1$ instead. Since $\mathcal{A}$ and $\mathcal{B}$ are linear expressions involving integers, these constraints are equivalent to the original ones. Finally, the constraints for which the inequality has the "$\geq$" direction can be handled easily using standard techniques.

The connection above established that Problem 4 can be solved in polynomial time when the number of input variables ℓ is fixed. To show that Problem 2 can be solved in polynomial time as well, we will do binary search on the interval $[0, \sum_{i=1}^{\ell} \alpha_i]$ where each time a feasible solution is obtained, it will be an improvement over the previously found solution. The termination condition is when the length of the interval becomes smaller than $1/b^2$. Since the algorithm for solving Problem 2 is invoked at most $\log b^2 \sum \alpha_i$ times, the binary search algorithm terminates in polynomial time. The details of the algorithm can be found in Algorithm 3. □

Input: $\alpha_i \in \mathbb{Q}, a, b \in \mathbb{N}^*$
Output: x
1: $\hat{c} := 0, \hat{d} := \sum \alpha_i$
2: **while** $\hat{d} - \hat{c} > \frac{1}{b^2}$ **do**
3: $t := \frac{\hat{c}+\hat{d}}{2}$
4: **if** Problem 4 has no solution for $c = t, d = \hat{d}$ **then**
5: $\hat{d} := t$
6: **else**
7: $\hat{c} := t$
8: **end if**
9: **end while**
10: return $x, k_1, \ldots, k_\ell$ computed for the last instance of Problem 4.

Algorithm 3: Algorithm for Problem 2 using an algorithm for Problem 4.

B Interesting Special Cases for Price Making

In the main text, we proved the approximation ratio guarantees of Mechanism ALL-OR-NOTHING, as a function of the market share. In this section, we will examine two interesting special cases, namely:

- When the budgets are *common*, i.e. it holds that $B_i = B$ for all buyers $i \in N$.
- The budgets are monotone in the valuations, i.e. $v_i \geq v_j \Leftrightarrow B_i \geq B_j$. We will call this market a *monotone market*.

Note that the second case is more general than the first, where for the right-hand side we have $B_i = B_j$ for all $i, j \in N$. We will prove that for those cases, Mechanism ALL-OR-NOTHING is optimal among all strategyproof mechanisms, for both the welfare and the revenue objective. For the welfare objective, the approximation ratio guarantee will be completely independent of the market share. For the revenue objective, the dependence will be rather weak; we prove that the bound holds in all markets except monopsonies. A monopsony is a market in which a single buyer can afford to buy all the items at a very high price.

Definition 3 *A market is a* monopsony*, if the buyer with the highest valuation v_1 has enough budget B_1 to buy all the units at a price equal to the second highest valuation v_2.*

Note that when the market is not a monopsony, that implies that the market share s^* is less than 1.[7]

Theorem 10 *When the market is monotone, the approximation ratio of Mechanism* All-or-Nothing *is*

- *at most 2 for the social welfare objective.*
- *at most 2 for the revenue objective when the market is not a monopsony.*

Furthermore, no strategyproof mechanism can achieve an approximation ratio smaller than $2 - \frac{4}{m+2}$ even in the case of common budgets.

Proof: First, note that the profile constructed in Theorem 9 is one where the budgets are common and therefore the lower bound extends to both cases mentioned above. Therefore, it suffices to prove the approximation ratio of Mechanism ALL-OR-NOTHING for both objectives, when the market is monotone.

We start from the social welfare objective and consider an arbitrary profile $\mathbf{v}$. Without loss of generality, we can assume that $\mathbf{v}$ is not trivial (otherwise the optimal allocation allocates 0 items in total) and note that the optimal envy-free price is $p^* = p_{min}$ and let $\mathbf{x}$ be the corresponding optimal allocation. Following the

[7] Note that instead of ruling out monopsonies, another approach would be to consider a different benchmark, that does not include the case of an omnipotent buyer, like the EFO$^{(2)}$ benchmark for revenue, see [25], Chapter 6.

arguments in the proof of Theorem 7, we establish that the according to $\mathbf{x}$ at most one additional semi-hungry buyer is allocated a positive number of units, compared to the allocation of Mechanism ALL-OR-NOTHING; let ℓ be that buyer and let x_ℓ be its optimal allocation.

The social welfare loss of Mechanism ALL-OR-NOTHING is $x_\ell \cdot v_\ell \leq v_\ell \cdot \lfloor B_\ell / v_\ell \rfloor$, i.e. the contribution of the the semi-hungry buyer that receives 0 items by ALL-OR-NOTHING, in contrast to the optimal allocation. Since the profile $\mathbf{v}$ is not trivial, there exists at least on other buyer j that receives $\min\{m, \lfloor B_j / v_\ell \rfloor\}$ units in the optimal allocation $\mathbf{x}$. If it receives m units, then $x_\ell = 0$ and the ratio on the profile is 1. Otherwise, the contribution to the welfare (for both the optimal allocation and the allocation of ALL-OR-NOTHING) from buyer j is $v_j \cdot \lfloor B_j / v_\ell \rfloor\} \geq v_j \cdot \lfloor B_\ell / v_\ell \rfloor\}$, since $v_\ell \leq v_j \Leftrightarrow B_\ell \leq B_j$ by the monotonicity of the market. Then we have:

$$\frac{\mathcal{SW}(AON)}{OPT} \geq \frac{OPT - v_\ell \cdot \lfloor B_\ell / v_\ell \rfloor}{OPT} = 1 - \frac{v_\ell \cdot \lfloor B_\ell / v_\ell \rfloor}{OPT} \geq 1 - \frac{v_\ell \cdot \lfloor B_\ell / v_\ell \rfloor}{(v_\ell + v_j) \cdot \lfloor B_\ell / v_\ell \rfloor} = 1 - \frac{v_\ell}{v_j + v_\ell} \geq \frac{1}{2}.$$

For the revenue objective, again let $p*$ be the optimal envy-free price and let $\mathbf{x}$ be the corresponding allocation. We consider two cases:

- $p^* = p_{min}$: The argument in this case is very similar to the one used above for the social welfare objective. In particular, since $p^* = p_{min} = v_\ell$, we now have that the loss in revenue from the semi-hungry buyer ℓ for Mechanism ALL-OR-NOTHING is at most $x_\ell \cdot v_\ell \leq v_\ell \cdot \lfloor B_\ell / v_\ell \rfloor$ whereas the contribution from buyer j is $v_\ell \cdot \lfloor B_j / v_\ell \rfloor\}$, which is at most $v_\ell \cdot \lfloor B_\ell / v_\ell \rfloor\}$ by the monotonicity of the market. Therefore, we have that:

$$\frac{\mathcal{REV}(AON)}{OPT} \geq \frac{OPT - v_\ell \cdot \lfloor B_\ell / v_\ell \rfloor}{OPT} = 1 - \frac{v_\ell \cdot \lfloor B_\ell / v_\ell \rfloor}{OPT} \geq 1 - \frac{v_\ell \cdot \lfloor B_\ell / v_\ell \rfloor}{2v_\ell \cdot \lfloor B_\ell / v_\ell \rfloor} = 1 - \frac{v_\ell}{2v_\ell} = \frac{1}{2}.$$

- $p^* > p_{min}$. In that case, the argument is exactly the same as in Case 2 of the proof of Theorem 7, which holds when the market share is less than 1, i.e. when the market is not a monopsony.

□

To complete the picture, we prove in the following that if the market is a monopsony, the approximation ratio of any strategyproof mechanism is unbounded. This can be captured by the following theorem.

Theorem 11 *If the market is a monopsony, the approximation ratio of any strategyproof mechanism for the revenue objective is at least $\mathcal{B}$ for any $\mathcal{B} > 1$, even if the budgets are public.*

Proof: Consider the following monopsony. Let $i_1 = \operatorname{argmax}_i v_i$, for $i = 1, \ldots, n$ be a single buyer with the highest valuation and denote $v_{i_1} = v_1$ for ease of notation. Similarly, let $i_2 \in \operatorname{argmax}_{i \in N \setminus \{i_1\}} v_i$ be

one buyer with the second largest valuation and let $v_{i_2} = v_2$. Furthermore, let $v_1 > v_i$ for all $i \neq i_1$ and $B_{i_1} = p \cdot m$, for some $v_2 < p \leq v_1$ i.e. buyer i_1 can afford to buy all the units at some price $p > v_2$. Additionally, let $B_{i_2} \geq v_2$, i.e. buyer i_2 can afford to buy at least one unit at price v_2.[8] Finally, for a given $\mathcal{B} > 1$ let v_2 and p be such that $\mathcal{B} = p/v_2$. Note that the revenue-maximizing envy-free price for the instance $\mathbf{v}$ is at least p and the maximum revenue is at least $p \cdot m$.

Assume for contradiction that there exists a strategyproof mechanism M with approximation ratio smaller than $\mathcal{B}$ and let p^* be the envy-free price output by M on $\mathbf{v}$. Since p^* is envy-free and $B_{i_1} > v_1 \cdot m$ and $B_{i_2} \geq v_2$, it can not be the case that $p^* < v$, otherwise there would be over-demand for the units. Furthermore, by assumption it can not be the case that $p^* = v_2$ as otherwise the ratio would be $\mathcal{B}$ and therefore it must hold that $p^* > v_2$.

Now let $\mathbf{v}'$ be the instance where all buyers have the same valuation as in $\mathbf{v}$ except for buyer i_1 that has value v_1' such that $v < v_1' < p^*$ and let $\tilde{p}$ be the envy-free price that M outputs on input $\mathbf{v}'$. If $\tilde{p} > v_1'$, then the ratio of M on the instance $\mathbf{v}'$ is infinite, a contradiction. If $\tilde{p} \leq v_1'$ and since $\tilde{p}$ is envy-free, it holds that $v_2 \leq \tilde{p} < p^*$. In that case however, on instance $\mathbf{v}$, buyer i_1 would have an incentive to misreport its valuation as v_1' and reduce the price. The buyer still receives all the units at a lower price and hence its utility increases as a result of the devation, contradicting the strategyproofness of M. □

C General Valuations

In the main text, we have considered the case of *linear* multi-unit markets. In this section, we study price taking buyers in *general multi-unit markets*, where the valuation of a buyer for different numbers of copies is not necessarily a linear function of that number. In the general case, each buyer has a valuation vector v_i, such that $v_{i,j}$ represents buyer i's valuation for receiving j units of the good.

$$u_i(p, x) = \begin{cases} v_{i,x_i} - p \cdot x_i, & \text{if } p \cdot x_i \leq B_i. \\ -\infty, & \text{otherwise}. \end{cases} \tag{1}$$

The *demand* of a buyer is defined as

$$D_i = \{y : y \in \operatorname*{argmax}_{x_i} u_i(p, x_i) \ \ , \ \ p \cdot y \leq B_i\},$$

i.e. a set consisting of allocations that make the buyer maximally satisfied at the chosen price and respect its budget constraint. The definition of an *envy-free pricing scheme* is the same as the one given in the Preliminaries section.

We will first prove that finding a social welfare-maximizing or a revenue-maximizing envy-free price and allocation is NP-hard by a reduction from SUBSET-SUM and then we prove that both problems admit

[8] Note that setting $B_{i_2} = B_{i_1}$ satisfies this constraint and creates a market with identical budgets, so the proof goes through for that case as well.

an FPTAS by using a variant of KNAPSACK called MULTI-CHOICE KNAPSACK. Note that since buyers have to specify their values $v_{i,j}$ for all possible number of units that they might receive, the input size to the problems here is n and m instead of n and $\log m$, which was the input size for the linear multi-unit market. This means that an NP-hardness result for the linear case would not imply NP-hardness for the general case and an FPTAS for the general case does not imply an FPTAS for the linear case.

Theorem 12 *Given a multi-unit market with general valuations, it is NP-hard to compute a revenue-maximizing or social welfare-maximizing envy-free price and allocation.*

Proof: We construct a reduction from the NP-complete problem SUBSET-SUM, which is defined next:

> *Given a universe of positive integers $\mathcal{U} = \{s_1, \ldots, s_n\}$ and an integer K, determine whether there exists a subset $S \subseteq U$ that sums up to exactly K.*

Given an input $\langle \mathcal{U}, K \rangle$ to SUBSET-SUM, we construct a market as follows. Let $N = \{1, \ldots, n\}$ be the set of buyers and $m = K$ the number of units. Set the budget of each buyer i to $b_i = s_i$ and the value of i for j units of the good to:

$$v_i(j) = \begin{cases} s_i & \text{if } j \geq s_i \\ 0 & \text{if } j < s_i \end{cases}$$

We show that a revenue of at least K can be obtained via an envy-free pricing in this market if and only if the SUBSET-SUM problem has a solution.

Suppose there is a price p and allocation $\mathbf{x}$ such that $(p, \mathbf{x})$ is envy-free and the revenue attained is at least K. Since each buyer i gets an optimal bundle, we have: $u_i(x_i, p) = v_i(x_i) - p \cdot x_i \geq 0$. There are three cases:

1). $x_i = s_i$. Then $s_i - p \cdot s_i \geq 0$, and so $p \leq 1$.

2). $x_i < s_i$. Then $v_i(x_i) = 0$, so $u_i(x_i, p) = 0 - p \cdot x_i \geq 0$, which means that either $p = 0$ or $x_i = 0$. Since $p > 0$ at any positive revenue, it must be that $x_i = 0$.

3). $x_i > s_i$. Then $v_i(x_i) = s_i$, so $u_i(x_i, p) = s_i - p \cdot x_i \geq 0$, so $p \leq s_i / x_i < 1$. Since the revenue is at least K and there are exactly K units, it cannot be that $p < 1$, and so this case can never occur.

From cases $(1-3)$ it follows that each buyer i gets either s_i or zero units and that the price is at most 1. Since the revenue is at least K and attained from selling K units, we get that in fact $p = 1$. Let $\mathcal{I} = \{i \in N \mid x_i > 0\}$ be the set of buyers that get a non-zero allocation. Then $Rev(p, \mathbf{x}) = \sum_{i \in \mathcal{I}} x_i \cdot p = \sum_{i \in \mathcal{I}} s_i = K$, which implies that $\mathcal{I}$ is a solution to the SUBSET-SUM problem.

For the other direction, let $\mathcal{I}$ be a solution to SUBSET-SUM. Set the price to $p = 1$ and the allocation $\mathbf{x}$ to $x_i = s_i$. It can be checked that $(p, \mathbf{x})$ is an envy-free pricing with revenue exactly K.

The reduction for social welfare is almost identical and exploits the fact that the equilibrium price achieving the target social welfare must, again, be equal to 1 (if it exists). □

Theorem 13 *Given a multi-unit market with general valuations, there exists an FPTAS for the problem of computing a welfare or revenue maximizing envy-free price.*

Proof: (sketch) The main idea is that for general valuations, the hardness no longer comes from guessing the price—in fact we can freely iterate over the set of candidate prices defined in Section 3.1 and the optimal solution is still guaranteed in this set for the same reason as that in the case of general valuations. Rather, the hardness comes from selecting the set of buyers to be allocated at a given price.

Thus if we had a black box that produced an approximate solution efficiently for each fixed price, then the problem would be solved. This black box will be the FPTAS for the 0-1 MULTI-CHOICE KNAPSACK PROBLEM; there are several such algorithms, such as the one due to Lawler [26], which for n items and precision $\epsilon > 0$, runs in time $O(n \log 1/\epsilon + 1/\epsilon^4)$.

The 0-1 MULTI-CHOICE KNAPSACK PROBLEM is as follows:

> We are given m classes C_1, C_2, . . . , C_m of items to pack in some knapsack of capacity C. Each item $j \in C_i$ has a value $p_{i,j}$ and a weight $w_{i,j}$. The problem is to choose at most one item from each class such that the value sum is maximized without the weight sum exceeding C.

Then our algorithm for maximizing revenue or welfare in a market with general valuations is as follows. For each price $p \in \mathcal{P}$, where $\mathcal{P}$ is the set of candidate prices, construct a 0-1 MULTI-CHOICE KNAPSACK instance, such that the ith category contains an item for every quantity (i.e. # of units) that can be sold to player i. That is, we create an element in category C_i for every possible integer $y \in \{0, \ldots, m\}$ with the property that player i's utility is non-negative should he be allocated exactly y units. Note that some numbers of units will be missing from C_i, exactly at those sizes at which i's utility would be negative.

- For the objective of maximizing revenue, set the weight of the item corresponding to y equal to y and its value equal to $p \cdot y$.
- For the objective of maximizing social welfare, set the weight of each item corresponding to y equal to y and its value to $v_{i,y}$.

Then set the total volume constraint to m and the total value (target k).

Then it can be seen that a value of K can be obtained in the 0-1 MULTI-CHOICE KNAPSACK problem if and only if a target revenue (or welfare respectively) of K can be obtained in the market. Note that if a knapsack solution omits taking an element from some class C_i, we can obtain an equivalent solution that is

still feasible where the item with zero value and zero weight from class C_i is included on top of all the other elements that are part of the solution.

Since we have an FPTAS for solving the 0-1 MULTI-CHOICE KNAPSACK instance, we can just this algorithm as a subroutine $n \times m$ times, which makes the total runtime still polynomial in the market size and $1/\epsilon$. □

D NP-hardness of Problem 1

A natural question would be whether one could use Problem 1 to construct a polynomial-time algorithm for finding a revenue-maximizing envy-free pricing scheme in linear multi-unit markets, for any number of buyers. In this subsection, we prove that unless P=NP, this is not possible. Note that the hardness of the problem does not imply NP-hardness of finding a revenue-maximizing envy-free pricing scheme in the market.

Theorem 14 *Problem 1 is NP-hard.*

To prove Theorem 14, we will construct a series of problems that each could be solved using a polynomial time algorithm for the previous problem, with the last problem being the well-known k-CLIQUE problem. Consider the following problem.

Problem 4 *Input:* $p_i, q_i \in \mathbb{N}^*$ *for* $i = 1, 2, \ldots, n$ *and* $a, b \in \mathbb{N}^*$

Output: $\arg\min \sum_{i=1}^{n} \frac{p_i x \mod q_i}{q_i x}$, $x \in [a, b] \cap \mathbb{N}^*$

It is not hard to see that Problem 1 is equivalent to Problem 4, and it is sufficient to consider the case when $x \in (k\Pi q_i, (k+1)\Pi q_i), k \in \mathbf{Z}$, otherwise it is trivial to obtain an objective function of value 0. Next, we will prove that the objective function satisfies a monotonicity condition when a is sufficiently large.

Lemma 13 *Denote* $r(x) = \sum \frac{p_i x \mod q_i}{q_i}$. *If* $a > n(\Pi q_i)^2$, *then*

$$r(x_1) \leq r(x_2) \Leftrightarrow \frac{r(x_1)}{x_1} \leq \frac{r(x_2)}{x_2}.$$

Proof: First, notice that following the discussion above about the interval in which we should be searching for x, it holds that $r(x_2) \geq r(x_1) + \frac{1}{\Pi q_i}$. This implies that

$$r(x_1)x_2 - r(x_2)x_1 \leq r(x_1)x_2 - r(x_1)x_1 - \frac{x_1}{\Pi q_i}.$$

If $r(x_1) \leq r(x_2)$, then $r(x_1)x_2 - r(x_2)x_1$ is at most $n\Pi q_i - \frac{x_1}{\Pi q_i}$ which is at most 0, since $x_2 - x_1 < \Pi q_i$, $r(x) \leq n$ for any x and $x_1 \geq a \geq n(\Pi q_i)^2$.

The other direction can be shown similarly. □

Now consider the following problem and observe that by Lemma 13, it is special case of Problem 4.

Problem 5 *Input:* $p_i, q_i \in \mathbb{N}^*$ *for* $i = 1, 2, \ldots, n$ *and* $a, b \in (n\Pi q_i, (n+1)\Pi q_i)$.

Output: $\arg\min \sum_{i=1}^{n} \frac{p_i x \mod q_i}{q_i}$, $x \in [a, b] \cap \mathbb{N}^*$

By subtracting $n\Pi q_i$ from the solutions to Problem 5, we obtain the following equivalent problem.

Problem 6 *Input:* $p_i, q_i \in \mathbb{N}^*$ *for* $i = 1, 2 \ldots, n)$ *and* $a, b \in (0, \Pi q_i) \cap \mathbb{N}^*$

Output: $\arg\min \sum_{i=1}^{n} \frac{p_i x \mod q_i}{q_i}$, $x \in [a, b] \cap \mathbb{N}^*$

As a special case of Problem 5, we double n and set $p_i = 1$ for n terms and $p_i = q_i - 1$ for the other n terms. We obtain the following problem.

Problem 7 *Input:* $a, b \in (0, \Pi q_i) \cap \mathbb{N}^*, q_i \in \mathbb{N}^*$ *for* $i = 1, 2, \ldots, n$

Output: $\arg\min \sum_{i=1}^{n} \frac{x \mod q_i}{q_i} + \frac{(q_i - 1)x \mod q_i}{q_i}, x \in [a, b] \cap \mathbb{N}^*$

Note that for each term in the sum of Problem 7, the term it equals to 0 if $q_i | x$ the term equals to 1 if $q_i \nmid x$. To minimize the objective function in Problem 7 we want to maximize the number of zero terms. Let $U = \{q_1, q_2, ..., q_n\}$ denote the set of the q_i and consider the following problem, which can be solved by solving Problem 7.

Problem 8 *Input:* $a, b \in (0, \Pi q_i) \cap \mathbb{N}^*, U$

Output: $\max k$ *such that* $X \subseteq U, |X| = k, q_i | x \; \forall q_i \in X, x \in [a, b] \cap \mathbb{N}^*$

Next we construct a special case of Problem 7. We first construct a base set W and then set each q_i to be the product of some elements in W. Let

$$W = \{M, M+1, M+2, ..., M+n^2-1\}, M > 2^n n^4.$$

We prove the following lemma.

Lemma 14 *If* $1 \leq x \leq (M + n^2 - 1)^k$, *then for any* $k (k < n^2)$ *different elements in* W *there exists* x *s.t.* x *is divisible by them simultaneously. There does not exist* x *s.t.* x *is divisible by any* $k+1$ *different elements in* W *simultaneously.*

Proof: The first part is obvious. For the second part, note that the greatest common divisor between any two elements in W is at most n^2, so if x is divisible by $k+1$ different elements in W simultaneously, $x > \frac{M^{k+1}}{\frac{1}{2}k(k+1)n^2}$, the right side is larger than $(M+n^2-1)^k$. □

For Problem 7, given $T < n^2$, we set $a = 1, b = (M+n^2-1)^T$, and each p_i to be the product of a subset Q_i of W. For this to be a special case of Problem 8, it has to be that $\Pi q_i > b$. To ensure this, we first consider the case when $T \geq \sum_i |Q_i|$, where the solution can easily be seen to be n. For the case when $T < \sum_i |Q_i|$, setting $b = (M+n^2-1)^T$ means that $b < \Pi q_i$ since the product of any $T+1$ elements in W is larger than $(M+n^2-1)^T$. Now x is divisible by at most T different elements in W simultaneously and we obtain the following problem as a special case.

Problem 9 *Input:* $Q_1, Q_2, ..., Q_n \subseteq W, T$

Output: $\max k$ *such that* $X \subseteq W, |X| \leq T, \sum_{i=1}^{n} I(Q_i \subseteq X) = k$

Now consider a graph $G = (V, E)$ where $|V| = n$, and let $Q_i = \{w_{i1}, w_{i2}, ..., w_{in}\}$, where $w_{ij} \in W$ and let $w_{ij} = w_{ji}$ if $(i,j) \in E$, $w_{ij} \neq w_{ji}$ if $(i,j) \notin E$. Since $|W| = n^2$ there are enough elements to construct Q_i. Given $k < n$, we set $T = nk - \frac{1}{2}k(k-1)$ and we obtain a special case of Problem 9.

Problem 10 *On input* k*, decide whether there exists* $X \in W$ *s.t.* $|X| \leq nk - \frac{1}{2}k(k-1), \sum_{i=1}^{n} I(Q_i \subseteq X) = k$

Finally, we will prove that Problem 10 is NP-hard, by a reduction to k-CLIQUE.

Lemma 15 *The answer to Problem 10 is yes if and only if Graph* G *has a k-clique.*

Proof: Note that the cardinality of the union of the k subsets Q_i equals to nk minus the total number of edges in the Q_i's corresponding subgraph of G. It is larger than $nk - \frac{1}{2}k(k-1)$ if the corresponding subgraph is not a clique, so X can not cover them. □

Since every problem was a special case of the previous one, this establishes Theorem 14.